

09/875136



## REDUCED AREA OF CROSSBAR AND METHOD OF OPERATION

## TECHNICAL FIELD OF THE INVENTION

This invention relates generally to multi-processor systems and more particularly to such systems and methods where the several processors are interconnectable to many  
5 different memory addressing spaces by a multi-port switch.

All of the following patent applications are cross-referenced to one another, and all have been assigned to Texas Instruments Incorporated. These applications have been concurrently filed and are hereby incorporated in this patent application by reference.

Attorney	<u>Docket</u>	<u>Title</u>
10	TI-14608	Multi-Processor With Crossbar Link of Processors and Memories and Method of Operation
15	TI-14655	SIMD/MIMD Reconfigurable Multi-Processor and Method of Operation
	TI-14656	Reconfigurable Communications for Multi-Processor and Method of Operation
20	TI-14657	Reduced Area of Crossbar and Method of Operation
	TI-14658	Synchronized MIMD Multi-Processors, System and Method of Operation
25	TI-14659	Sliced Addressing Multi-Processor and Method of Operation
30	TI-14660	Ones Counting Circuit and Method of Operation
	TI-14661	Memory Circuit Reconfigurable as Data Memory or Instruction Cache and Method of Operation
35	TI-14662	Imaging Computer and Method of Operation
	TI-14704	Switch Matrix Having Integrated Crosspoint Logic and Method of Operation
40		

### BACKGROUND OF THE INVENTION

In the world of computers and processors there is an unrelenting drive for additional computing power and faster calculation times. In this context, then, systems in which  
5 several processors can be combined to work in parallel with one another are necessary.

Imaging systems which obtain visual images and perform various manipulations with respect to the data and then control the display of the imaged and stored data  
10 inherently require large amounts of computations and memory. Such imaging systems are prime candidates for multi-processing where different processors perform different tasks concurrently in parallel. These processors can be working together in the single instruction, multiple  
15 data mode (SIMD) where all of the processors are operating from the same instruction stream but obtaining data from various sources or the processors can be working together in the multiple instruction, multiple data mode (MIMD) where each processor is working from a different set of  
20 instructions and working on data from different sources. For different operations, different configurations are necessary.

In a multi-processor system each processor can have several buses or ports for the communication of data.  
25 Thus, assuming two buses for data and one bus for instructions, and assuming only four processors in the system, a minimum of twelve buses must be switched. When it is realized that additional buses may be required for master processors and control processors to handle  
30 simultaneous data input/output on a particular memory module and processing via a particular processor on other memory modules, the problem is compounded. In some situations it may be desirable to isolate certain memories for access only by a particular processor, such as a master  
35 processor.

Making the problem even more severe is the fact that in a multi-processing system the true power comes from the ability of any processor to communicate with any memory at any time combined with the ability of the processors to  
5 communicate with each other, all occurring simultaneously.

There is thus a need in the art for a system which handles multi-processors having multi-memories such that the address space from all of the memories is available to one or more processors concurrently even when the  
10 processors are handling different instruction streams.

One method of solving the huge interconnection problem in complex systems such as the image processing system shown in one embodiment of the invention is to construct the entire processor as a single device. Conceptually this  
15 might appear easy to achieve, but in reality the problems are complicated.

First of all, an architecture must be created which allows for the efficient movement of information, while at the same time consuming a minimum amount of precious silicon  
20 chip space in order to achieve a high performance to cost ratio. The architecture must allow a very high degree of flexibility, since once fabricated, it cannot easily be modified for different applications. Also, since the processing capability of the system will be high, there is  
25 a need for high bandwidth of each data input/output signal which moves information on and off the chip. This is so since the physical number of leads which can attach to any one chip is limited.

It is also desirable to design an entire parallel  
30 processor system, such as an image processor, on a single silicon chip while maintaining the system flexible enough to satisfy wide ranging and constantly changing operational criteria.

It is further desirable to construct such a single  
35 chip parallel processor system where the processor memory

0935136-050501  
TOP SECRET



[illegible]

SUMMARY OF THE INVENTION

These problems have been solved by designing a multi-processing system to handle image processing and graphics and by constructing a crossbar switch capable of interconnecting any processor with any memory in many configurations for the interchange of data. The system is capable of connecting  $n$  parallel processors to  $m$  memories where  $m$  is greater than  $n$ . The system, in one embodiment, has four processors capable of operating in either the SIMD or MIMD modes. Each processor has three buses, two for data and one for instructions. The data ports are divided into global and local ports. The global port of each processor is arranged to access, via a crossbar switch, any one of the individual addressable memory spaces. The local port is arranged to access, via the same crossbar switch, only a subset of the addressable memory spaces, while the instruction port is even more limited in that it can access only one memory. The limitations imposed on the local and instruction ports allow for a minimization of the crossbar buses, thereby saving substrate space.

The crossbar switch allows the processors to be tied together on a cycle by cycle basis for the purposes of allowing a common loading of data or instructions from memory.

Thus, it is a technical advancement that a crossbar switch has been arranged to allow several processors to access several memories on a random basis and to do so concurrently on a cycle by cycle basis. The processors may still communicate with one another and with a common memory during any cycle while communicating with separate memories during other cycles.

The problems inherent with constructing a single chip image processor having a high degree of versatility have been solved by the architecture of establishing a multi-link, multi-bus crossbar switch between the individual

5

10

### BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and for further advantages thereof, reference is now made to the following detailed description taken in  
5 conjunction with accompanying drawings in which

FIGURES 1 and 2 show an overall view of the elements of the image processing system;

FIGURE 3 shows a series of image processing systems interconnected together into an expanded system;

10 FIGURE 4 shows details of the crossbar switch matrix interconnecting the parallel processors and the memories;

FIGURES 5 and 6 show prior art configurations;

FIGURE 7 shows an improved configuration;

15 FIGURES 8 and 9 show prior art schematic representations of processor memory interaction;

FIGURE 10 shows some reconfigurable modes of operations of an improved multi-processor;

FIGURE 11 is a graph showing some algorithms and control for the image processing system;

20 FIGURES 12-15 show image pixel flow for SIMD and MIMD operational modes;

FIGURE 16 shows the interrupt polling communication between the processors;

25 FIGURE 17 shows a schematic representation of the layout of the processors and memory interconnected by the crossbar switch;

FIGURES 18 and 19 show details of the crosspoints of the crossbar switch;

30 FIGURE 20 is a graph of wave forms of the contention logic for memory access;

FIGURES 21-23 show the synchronization control between processors;

FIGURES 24-27 show details of the sliced addressing technique;

FIGURE 28 shows details of the rearrangement of the instruction data memory for the SIMD/MIMD operational modes;

FIGURE 29 shows details of a master processor;

5        FIGURES 30-34 show details of the parallel processors;

FIGURES 35-45 show figures useful in understanding methods of operation of the parallel processor;

FIGURES 46-48 show an image processor operating as a personal computer;

10       FIGURES 49-52 show system arrangements for use of the imaging system on a local and remote basis;

FIGURE 53 is a functional block diagram of an imaging system;

15       FIGURE 54 is a logic schematic of the ones counting circuit matrix;

FIGURE 55 is a logic schematic of a minimized matrix of the ones counting circuit;

FIGURE 56 is an example of an application of a ones counting circuit;

20       FIGURE 57 shows a block diagram of the transfer processor;

FIGURE 58 shows a block diagram of the parallel processor system used with a VRAM; and

25       FIGURES 59-64 show various operational mode relationships.

TOP SECRET

### DETAILED DESCRIPTION OF THE INVENTION

Prior to beginning a discussion of the operation of the system, it may be helpful to understand how parallel processing systems have operated in the prior art.

5       FIGURE 5 shows a system having parallel processors 50-53 accessing a single memory 55. The system shown in FIGURE 5 is typically called a shared memory system where all of the parallel processors 50-53 share data in and out of the same memory 55.

10       FIGURE 6 shows another prior art system where memory 65-68 is distributed with respect to processors 60-63 on a one-for-one basis. In this type of system, the various processors access their respective memory in parallel and thus operate without memory contention between the  
15 processors. The system operating structures shown in FIGURES 5 and 6, as will be discussed hereinafter, are suitable for a particular type of problem, and each is optimized for that type of problem. In the past, systems tended to be either shared or distributed.

20       As processing requirements become more complex and the speed of operation becomes critical, it is important for systems to be able to handle a wide range of operations, some of which are best performed in the shared memory mode, and some of which are best performed in a distributed  
25 memory mode. The structure shown in FIGURES 1 and 2 accomplishes this result by allowing a system to have parallel processing working both in the shared and in the distributed mode. While in these modes, various operational arrangements such as SIMD and MIMD can be  
30 achieved.

### Multi-Processors and Memory Interconnection

As shown in FIGURE 1, there is a set of parallel processors 100-103 and a master processor 12 connected to  
35 a series of memories 10 via a cycle-rate local connection

0987436-060604

network switch matrix 20 called a crossbar switch. The crossbar switch, as will be shown, is operative on a cycle by cycle basis to interconnect the various processors with the various memories so that different combinations of distributed and shared memory arrangements can be achieved from time to time as necessary for the particular operation. Also, as will be shown, certain groups of processors can be operating in a distributed mode with respect to certain memories, while other processors concurrently can be operating in the shared mode with respect to each other and with respect to a particular memory.

Another view of the system is shown in FIGURE 2 in which the four parallel processors 100, 101, 102, 103 are shown connected to memory 10 via switch matrix 20 which is shown in FIGURE 2 as a distributed bus. Also connected to memory 10 via crossbar switch 20 is transfer processor 11 and master processor 12. Master processor 12 is also connected to data cache 13 via bus 171 and instruction cache 14 via bus 172. The parallel processors 100 through 103 are interconnected via communication bus 40 so that the processors, as will be discussed hereinafter, can communicate with each other and with master processor 12 and with transfer processor 11. Transfer processor 11 communicates with external memory 15 via bus 21.

Also in FIGURE 2, frame controllers 170 are shown communicating with transfer processor 11 via bus 110. Frame controllers 170 serve to control image inputs and outputs as will be discussed hereinafter. These inputs can be, for example, a video camera, and the output can be, for example, a data display. Any other type of image input or image output could also be utilized in the manner to be more fully discussed hereinafter.

Crossbar switch 20 is shown distributed, and in this form tends to mitigate communication bottlenecks so that

communications can flow easily between the various parts of the system. The crossbar switch is integrated on a single chip with the processors and with the memory thereby further enhancing communications among the system elements.

5 Also, it should be noted that fabrication on a chip is in layers and the switch matrix may have elements on various different layers. When representing the switch pictorially, it is shown in crossbar fashion with horizontals and verticals. In actual practice these may be  
10 all running in the same direction only separated spatially from one another. Thus, the terms horizontal and vertical, when applied to the links of the switch matrix, may be interchanged with each other and refer to spatially separated lines in the same or different parallel planes.

15 Digressing momentarily, the system can operate in several operational modes, one of these modes being a single instruction multiple data (SIMD) mode where a single instruction stream is supplied to more than one parallel processor, and each processor can access the same memory or  
20 different memories to operate on the data. The second operational mode is the multiple instruction, multiple data mode (MIMD) where multiple instructions coming from perhaps different memories operate multiple processors operating on data which comes from the same or different memory data  
25 banks. These two operational modes are but two of many different operational modes that the system can operate in, and as will be seen, the system can easily switch between operational modes periodically when necessary to operate the different algorithms of the different instruction  
30 streams.

Returning briefly to FIGURE 1, master processor 12 is shown connected to the memories via crossbar switch 20. Transfer processor 11, which is also shown connected to crossbar switch 20, is shown connected via bus 21 to  
35 external memory 15. Also note that as part of memory 10,



there are several independent memories and a parameter memory which will be used in conjunction with processor interconnection bus 40 in a manner to be more fully detailed hereinafter. While FIGURE 2 shows a single  
 5 parameter memory, in actuality the parameter memory can be several RAMS one per processor which makes communication more efficient and allows the processors to communicate with the RAMS concurrently.

FIGURE 4 shows a more detailed view of FIGURES 1 and  
 10 2 where the four parallel processors 100-103 are shown interconnected by communication bus 40 and also shown connected to memory 10 via crossbar switch matrix 20. The various crosspoints of the crossbar switch will be referred to by their coordinate locations starting in a lower left  
 15 corner with 0-0. In the numbering scheme, the vertical number will be used first. Thus, the lower left corner crosspoint is known as 0-0, and the one immediately to the right in the bottom row would be 1-0. FIGURE 19 which will be discussed hereinafter, shows the details of a particular  
 20 crosspoint, such as crosspoint 1-5. Continuing now in FIGURE 4, the individual parallel processors, such as parallel processor 103, are shown having a global data connection (G), a local data connection (L) and an instruction connection (I). Each of these will be detailed  
 25 hereinafter, and each serves a different purpose. For example, the global connection allows processor 103 to be connected to any of the several individual memories of memory 10, which can be for data from any of the various individual memories.

30 The local memory ports of the parallel processors can each address only the memories that are served by three of the vertical switch matrix links immediately opposite the processors. Thus, processor 103 can use verticals 0, 1 and 2 of crossbar 20 to access memories 10-16, 10-15 and 10-14  
 35 for data transfer in the MIMD mode. In addition, while in

T09090" 060601

the MIMD mode, memory 10-13 supplies an instruction stream to processor 103. As will be seen, in SIMD mode all of the instructions for the processors come from memory 10-1. Thus, instruction memory 10-13 is available for data. In  
 5 this situation, the switch is reconfigured to allow access via vertical 4 of crossbar 20. The manner in which crossbar 20 is reconfigured will be discussed hereinafter.

As shown in FIGURE 4, each parallel processor 100-103 has a particular global bus and a particular local bus to  
 10 allow the processor access to the various memories. Thus, parallel processor 100 has a global bus which is horizontal 2 of crossbar 20, while parallel processor 101 has a global bus which is horizontal 3 of crossbar 20. Parallel processor 102 has as its global bus horizontal 4, while  
 15 parallel processor 103 has as its global bus horizontal 5.

The local buses from all of the processors share the same horizontal 6. However, horizontal 6, as can be seen, is separated into four portions via three-state buffers 404, 405 and 406. This effectively provides isolation on  
 20 horizontal 6 so that each local input to each processor can access different memories. This arrangement has been constructed for efficiency of layout area on the silicon chip. These buffers allow the various portions to be connected together when desired in the manner to be  
 25 detailed hereinafter for the common communication of data between the processors. This structure allows data from memories 10-0, 10-2, 10-3 and 10-4 to be distributed to any of the processors 100-103.

When the processor is operating in the MIMD  
 30 operational mode, the instruction port of the processors, for example, the instruction port of processor 103, is connected through crosspoint 4-7 to instruction memory 10-13. In this mode crosspoints 4-2, 4-3, 4-4, 4-5 and 4-6, as well as 4-1, are disabled. In this mode crosspoint 4-0  
 35 is a dynamically operative crosspoint, thereby allowing the

transfer processor to also access instruction memory 10-13, if necessary. This same procedure is available with respect to crosspoint 9-7 (processor 102) and crosspoint 14-7 (processor 101).

5 When the system is in the SIMD mode crosspoint 4-7 is inactive, and crosspoints 4-2 through 4-6 may be activated, thereby allowing memory 10-13 to become available for data to all of the processors 100-103 via vertical 4 of crossbar 20. Concurrently, while in the SIMD mode buffers 401, 402  
10 and 403 are activated, thereby allowing instruction memory 10-1 to be accessed by all of the processors 100-103 via their respective instruction inputs. If buffer 401 is activated, but not buffers 402 and 403, then processors 100 and 101 can share instructor memory 10-1 and operate in the  
15 SIMD mode while processors 102 and 103 are free to run in MIMD mode out of memories 10-13 and 10-9 respectively.

Crosspoints 18-0, 13-0, 8-0 and 3-0 are used to allow transfer processor 11 to be connected to the instruction inputs of any of the parallel processors. This  
20 communication can be for various purposes, including allowing the transfer processor to have access to the parallel processors in situations where there are cache misses.

FIGURE 7 is a stylized diagram showing the operation  
25 of parallel processors 100-103 operating with respect to memories 55 and 55A in the shared mode (as previously discussed with respect to FIGURE 5) and operating with respect to memories 65-68 in the distributed mode (as previously discussed with respect to FIGURE 6). The manner  
30 of achieving this flexible arrangement of parallel processors will be discussed and shown to depend upon the operation of crossbar switch 20 which is arranged with a plurality of links to be individually operated at crosspoints thereof to effect the different arrangements  
35 desired.

FOR "SECRET" 050601

Before progressing to discuss the operation of the crossbar switch, it might be helpful to review FIGURE 3 and alternate arrangements where a bus 34 can be established connected to a series of processors 30-32, each processor 5 having the configuration shown with respect to FIGURES 1 and 2. External memory 35 is shown as a single memory 15, the same memory discussed previously. This memory could be a series of individual memories, both local and located remotely. The structure shown in FIGURE 3 can be used to 10 integrate any number of different type of processors together with the image system processor discussed herein, assuming that all of the processors access a single global memory space having a unified addressing capability. This arrangement also assumes a unified contention arrangement 15 for the memory access via bus 34 so that all of the processors can communicate and can maintain order while they each perform their own independent operations. Host processor 33 can share some of the policing problems between the various processors 30-32 to assure an orderly 20 flow of data via bus 34.

#### Image Processing

In image processing there are several levels of operations that can be performed on an image. These can be 25 thought about as being different levels with the lowest level being simply to message the data to perform basic operations without understanding the contents of the data. This can be, for example, removal of extraneous specks from an image. A higher level would be to operate on a 30 particular portion of the data, for example, recognizing that some portion of the data represents a circle, but not fully understanding that the circle is one part of a human face. A still higher operational aspect of image processing would be to process the image understanding that 35 the various circles and other shapes form a human image, or

other image, and to then utilize this information in various ways.

Each of these levels of image processing is performed most efficiently with the processors operating in a particular type of operational mode. Thus, when operations are performed on data locally grouped together without an attempt to understand the entire image, it is usually more efficient to use the SIMD operational mode where all, or a group of, processors operate from a single instruction and from multiple data sources. When operating in a higher mode where image pixel data is required from various aspects of the entire image in order to understand the entire image, the most efficient operational mode would be the MIMD mode where the processors each operate from individual instructions.

It is important to understand that when the system is operating in the SIMD mode, the entire pixel image can be processed through the various processors operating from a single instruction stream. This would be, for example, when the entire image is to be cleaned, or the image is enhanced to show various corners or edges. Then all of the image data passes through the processors in the SIMD mode, but at any one time data from various different areas of the image cannot be processed in a different manner for different purposes. The general operational characteristic of a SIMD operation is that at any period of time a relatively small amount of the data with respect to the entire image is being operated on. This is followed, in sequential fashion, by more data being operated on in the same manner.

This is in contrast to the MIMD mode where data from various parts of the image is being processed concurrently, some using different algorithms. In this arrangement, different instructions are operating on different data at the same time to achieve a desired result. A simple

example would include many different SIMD algorithms (like clean, enhance, extract) operating concurrently or pipelined on many different processors. Another example with MIMD would include the implementation of algorithms with the same data flow although using unique arithmetic or logical functions.

FIGURES 8 and 9 show the prior art form of the SIMD and MIMD processors with their respective memories. These are the preferred typologies for SIMD/MIMD for image processing. The operational modes of the system will be discussed more fully with respect to FIGURES 59-64. In general, data path 80 of FIGURE 8 corresponds to data path 6010, 6020, 6030 and 6040 of FIGURE 60, while processor 90 of FIGURE 9 corresponds to processor 5901, 5911, 5921, 5931 of FIGURE 59. The controller (6002 of FIGURE 60) for the data paths is not shown in FIGURE 8.

#### Reconfigurable SIMD/MIMD

FIGURE 10 shows the reconfigurable SIMD/MIMD topology of this invention where several parallel processors can be interconnected via crossbar switch 20 to a series of memories 10 and can be connected via a transfer processor 11 to external memory 15, all on a cycle by cycle basis.

One of the problems of operating in the MIMD topology is that data access can require high bandwidth as compared to operation in the SIMD mode where the effective data flow is on a serial basis or is emulated in the topology. Thus, in the SIMD mode, the data typically flows sequentially through the various processors from one processor to the next. This can be a blessing as well as a problem. The problem arises in that all of the data of the image has to be processed in order to arrive at a certain point in the processing. This is accomplished in the SIMD mode in a serial fashion. However, the MIMD mode solves this type of a problem because data from the individual memories can be

obtained at any time in the cycle, as contrasted to the operation in the SIMD where the shared memory can only be accessed upon a serial basis as the data arrives.

However, the MIMD mode has operational bottlenecks  
 5 when it is required to have interprocessor communication since then one processor must write the data to a memory and then the other processor must know the information is there and then access that memory. This can require several cycles of operational time and thus large images  
 10 with vast pixel data could require high processing times. This is a major difficulty. In the structure of FIGURE 10, as discussed, these problems have been overcome because the crossbar switch can serve to, on a cycle by cycle basis if necessary, interconnect various processors together to work  
 15 from a single instruction for a period of time or to work independently so that data which is stored in a first memory can remain in that memory while a different processor is for, one cycle or for a period of time, connected to that same memory. In essence, in some of the  
 20 prior art, the data must be moved from memory to memory for access by the various processors, which in the instant system the data can remain constant in the memory while the processors are switched as necessary between the memories. This allows for complete flexibility of processor and  
 25 memory operation as well as optimal use of data transfer resources.

A specific example of the processing of data in the various SIMD and MIMD modes can be shown with respect to FIGURES 12 and 13. In FIGURE 12 there is shown an image  
 30 125 having a series of pixels 0-n. Note that while in the image a row is shown having only four pixels, this is by way of example only, and a typical image would have perhaps a thousand rows, each row having a thousand pixels. At any one point in time the number of pixels in a row and the  
 35 number of rows will vary. For our purposes, we will assume

FIGURE 10

5

10

25

30

35



memory locations. In this mode the MIMD operation would be faster than the SIMD operation except in situations where data would have to move from processor 121 to processor 120, in which case there would have to be a movement of data in the memory bank. This interprocessor data movement could be required, for example, in situations where data processed from a particular region is important in determining how to process data from another region, or for determining exactly what the total image represents. Just as it is difficult to determine the shape of an elephant from a grasp of its trunk, it is equally difficult to obtain meaningful information from an image without access to different portions of the pixel data.

Turning now to FIGURE 14, there is graphically illustrated a system utilizing the present invention. Crossbar switch 20 allows processors 100-103 to access individual memories M1-M4 of memory 10, and on a cycle by cycle basis. The structure shown in FIGURE 14 allows the operation described in FIGURE 12 with respect to the SIMD operation such that the data in the memory elements, M1-M4 remains stationary and the connections from the processor switch. The continual flow of the process is enhanced by having more memory elements than actually utilized by the processors at a given instance. Thus, data can move in and out from these "extra" memory elements, and these extra elements can be cycled into the operational stream. In such an arrangement, data in and data out memory elements would, on a cycle by cycle basis, be different memory elements. Note that the data in and data out memories are switched through the crossbar and thus can be positioned in any of the memory elements. Thus, instead of moving the data between memories, the processor connection is sequentially changed.

Turning now to FIGURE 15, the MIMD mode is shown such that processors 100-103 are connected through crossbar

switch 20 to various memories. Typically, these connections would last through several cycles and thus, the processors each would be connected to the respective memories for a period of time. While this is not  
 5 necessary, it would be the most typical operation in the MIMD mode. For any processor, or group of processors operating in the MIMD mode of FIGURE 15, crossbar switch 20 can, on a cycle by cycle basis, be operated so that data from a particular memory element is immediately made  
 10 available to any of the other processors so that the data can either be cycled through the other processors or operated on a one-time basis.

#### Reconfigurable Interprocessor Communication

15 FIGURE 16 shows the diagram of interprocessor communication when the system is operating in the MIMD mode when the various processors must communicate with each other. A processor, such as processor 100, sends a message through crossbar switch 20 to the shared parameter memory  
 20 while at the same time registering a message (interrupt) in the destination processor that a parameter message is waiting. The destination processor, which can be any one of the other processors such as processor 102, then via crossbar switch 20 accesses the shared parameter memory to  
 25 remove the message. The destination process, for example, then could reconfigure itself in accordance with the received message. This reconfiguration can be internal to provide a particular system mode of operation or can be an instruction as to which memories to access and which  
 30 memories not to access for a period of time.

The question of accessing memories (contention) is important because a processor can waste a lot of time trying to access a memory when another processor is using that memory for an extended period. The efficient  
 35 operation of the system would be very difficult to achieve

09875436-060604  
 109090-957860

without the interprocessor coupling via the communication link.

Another type of message which is communicated between the processors relates to the synchronization of the processors. These messages and the precise manner in which synchronization is accomplished will be discussed hereinafter. FIGURE 2 shows the full system arrangement where the processors are interconnected for interrupting or polling between them to control sync, memory and crossbar allocation on a cycle by cycle basis.

It is the communication links between the processors which function outside of the crossbar switch that supports a more efficient utilization of the memory. The number of cycles that are required to switch operational modes, for example between SIMD and MIMD, is dependent upon the amount of other operations which must be performed. These other operations are, for example, loading of code in various instruction memories and the loading of data into data memories for subsequent operation. The external communications help this function by establishing which memories a particular processor may access and instructing all of the processors as to their ability to access memories so that the processors are not waiting in line for access when the access is being denied.

The instructions between processors can be by interrupt and by polling. The interrupt can be in any one of the well-known interrupt configurations where data can be transmitted with a flag to point to particular message locations within the shared parameter memory or can operate directly on a pointer basis within the processor. The ability to establish on a cycle by cycle basis which processor has access to which memory is important in establishing the ability of the system to operate in the MIMD mode so that data can reside in a particular memory, and the processors which have access to that data are

continually shifted. Using this arrangement then, several cycles of time, which would be required to move data from memory to memory if the memories were on a fixed relationship to processors, are dramatically eliminated.

- 5 The communication link includes the master processor.

#### Transfer Processor

- Transfer processor 11 shown in FIGURES 1 and 2 and in FIGURE 57 transfers data between external memory and the various internal memory elements. Transfer processor 11 is designed to operate from packet requests such that any of the parallel processors or the master processor can ask transfer processor 11 to provide data for any particular pixel or a group of pixels or data, and the transfer processor will transfer the necessary data to or from external and internal memory without further processor intervention instructions. This then allows transfer processor 11 to work autonomously and to process data in and out of the system without monitoring by any of the processors. Transfer processor 11 is connected to all of the memories through switch matrix 20 and is arranged to contend with the various links for access to the memories. Transfer processor 11 for any particular link may be assigned the lowest priority and access a memory when another processor is not accessing that memory. The data that is being moved by the transfer processor is not only the data for processing pixels, but instruction streams for controlling the system. These instruction streams are loaded into the instruction memory via crossbar switch 20.
- 30 Transfer processor 11 can be arranged with a combination of hardware and software to effect the purpose of data transfer.

FOR SECRET

### Master Processor

The master processor, shown in more detail in FIGURE 29, is used for scheduling and control of the entire system, including the control of the transfer processor as well as the interaction between the various processors. The master processor has a connection through the crossbar switch to all of the memories and is interconnected with the other processors on the communication channel. The master processor can control the type of data and the manner in which the data is obtained by the transfer processor depending upon the pixel information and the particular purpose for which the information is being obtained. Thus, regions of the image can be scanned under different scan modes depending upon the purpose for the scan. This is controlled by the master processor working in conjunction with the parallel processors. The parallel processors may each also control the transfer processor, either alone or in conjunction with the master processor, again depending upon the purpose for the operation.

The contention for the memory to the crossbar switch can be arranged such that the parallel processors have higher priority, the master processor has lower priority, and the transfer processor has third or lowest priority for any particular memory on a particular link.

FIGURE 11 shows a listing of various operations or algorithms which the imaging processing system would typically perform. A typical type of operation would be optical character recognition, target recognition or movement recognition. In each of these situations, the associated image processing would be controlled by the kind of operations to be performed.

In FIGURE 11, the types of operations which are typically performed by the parallel processors are shown below line 1100 and the types of operations which are

typically performed by the master processor are shown above line 1100. While this arrangement of operations is arbitrarily divided between the master processor and the parallel processors, the types of operations required to  
5 achieve the various operations shown tend to make them more suitable for either the master processor or the parallel processor.

As an example of image processing starting from an image and working higher in the hierarchy of operations,  
10 the image is first received by image enhancement 1111. In some situations it is necessary to compress or decompress the image via boxes 1112 and 1113. The image is then moved upwards through the various possibilities for edge extraction 1109, line linkage 1107, corner or vertices  
15 recognition 1105, histogram 1110, statistical properties 1108 and segmentation 1106. These boxes can all be skipped and the image provided directly to template matching 1102 for the purpose of determining the image identification 1101. There are various methods of achieving this  
20 identification, all of which are not necessary for every image, and all of which are well known in the art as individual algorithms or methods.

Enhancement block 1111 is a process which essentially cleans an image, removes extraneous signals and enhances  
25 details of the image, such as lines. Box 1109, edge extraction, is a process which determines the causes or existence of edges in an image. Box 1107 connects all the lines which have been extracted from the image and links them together to form longer lines. The process then  
30 removes extraneous dashes caused by inconsistencies in the data. Box 1105, corners and vertices, is an algorithm which determines where the corners of an image might be located. Once these geometric shapes are found, a process of grouping and labeling, block 1104, can then be used to

09875436-060601

identify major groupings of objects, such as circles and rectangles.

At this point, the operations have centered their focus on a smaller region of the image whereas in block 5 1111 the entire image is typically operated on. An alternate path after every enhancement is to perform statistical analysis, such as a histogram, 1110, of the intensities of the pixels. One purpose of a histogram is to discover the number of ones or the number of ones in a 10 particular axis or projection which would then be useful statistical information to quantify the presence of some object or orientation of an object. This will be discussed hereinafter.

Block 1108, statistical properties, then extracts from 15 these histograms the proper statistical properties. Continuing upward, block 1106 is a process of segmentation whereby the statistical properties could be used to segment different objects. As an example, several disconnected objects could then be quite easily segmented. Then through 20 the progression to grouping and labeling 1104, where an image has different objects identified with specific labels. Connector component algorithms are typical in this area. At this point also certain geometric features can be analyzed, particularly the perimeter of the object. Other 25 shape descriptors, Euler numbers, and a description of the surface can be obtained and used for future matching operations. Matching operations level 1102 is reached where similar information which is stored as templates or libraries are accessed and compared against the data that 30 is extracted from the lower level. This can be either geometric, surface description or optical flow information. Once a match has occurred, these matches then are statistically weighted to determine the degree of certainty that an object has been identified as shown by block 1101. 35 Once we have identified objects, we will in some

applications such as stereopsis or motion have a three dimensional representation of the world knowing what the objects are and where they are placed in the world. At this point we can then re-render the scene using a graphics pipeline as shown by the right side of FIGURE 11.

The first block, geometric model 1114, identifies a representation of this scene which basically is three coordinates showing position and a geometric description of the object such as its shape, density and reflective properties. At this point, depending upon the type of object, several different routes would be used to render the scene. If there were simple characters, two dimensional transforms would be employed. If they were more complex, three dimensional worlds would be created. A hand waving in front of a computer for use as a gesture input device would use this method and implement function 1116, which is a three dimensional transform. This would transform the input into a new coordinate system, either by translating scaling or rotating the three dimensional coordinates via 3D transform block 1116. Certain objects would be occluded by other objects. Again in the hand example, some fingers may be occluded by other fingers, and this operation using visibility block 1117 would then ignore the parts that were not visible. As we move down in FIGURE 11 to shaded solid box 1118 we find a process which would generate gray scale or pixel information to give a smooth shaded solid image which would be more realistic and more lifelike than taking the other route down to clipping box 1120. Clipping box 1120 essentially clips things that are out of the field of view of the scene that is being generated.

In a special case of rendering fonts on a computer screen or on a laser printer or such, box 1119, font compilation, would be used to create sophisticated fonts of multiple sizes and shapes. Then the final process in the

09875436-060604  
PAGE 28



graphics program would be actually to draw the objects, via block 1121, which might be as simple as drawing dots and lines that connect the dots. We are now back at the original level of image enhancement 1111 and have recreated  
 5 a synthetic representation of original image based upon a model which has been derived from that original image.

It is understood that once a character is recognized or a movement is recognized, an output can be obtained, either in binary code or otherwise, to control further  
 10 processing of the same image via output control 1122 by the operation and the combination of the parallel processors and the master processor working with the image processing system.

Generally, the boxes shown below line 1100 are  
 15 typically operationally efficient to be performed in the SIMD mode and require a vast amount of processing. These are performed with the parallel processing operation. The operations above line 1100 require relatively less processing capabilities and are less bandwidth intensive.  
 20 Accordingly, they are performed by a single processor. Also note that with respect to the operations, as the hierarchy moves upwards on the chart the likelihood is that the MIMD operations would be the preferred operation. Often the SIMD and MIMD operations overlap, and both types  
 25 of operational modes are required.

The main reason why two different types of processors are necessary is because of the level of the processing. High level processing, as performed by the master processor, preferably uses floating point arithmetic for  
 30 high precision. High precision floating point processors require more real estate space and are slower to operate from non-floating point processors. Therefore, if all of the processors were the same, there could be fewer processors on a given chip which would increase the problem  
 35 of bandwidth and slow down the operation of the system.

09876543210 " 09876543210

On the other hand, the low level processors do not require floating point arithmetic and thus can be made faster and smaller, which in turn allows more processors to be constructed on a given chip. The bus structure shown  
5 utilizing a crossbar switch can therefore take several different types of processors as required and switch them into the system to perform portions of every operation if necessary.

The master processor is designed to operate primarily  
10 on lists such as information lists and display lists, whereas the parallel processors are intended to operate on arrays. At the low level image processing most of the information can be described as two dimensional arrays, whereas at the higher level, the information is described  
15 as lists of multidimensional coordinates. The manipulation of these two different types of data representations requires different processing structures which is another motivation for the master and parallel processors having different structures.

The master processor of the preferred embodiment would have features similar to a RISC processor which is primarily intended for general purpose computing operations, whereas the parallel processors are more like digital signal processors (DSP) which tend to be  
20 specialized processors for arithmetic operations. Thus, the system could be optimized for the types of information processing required for image systems, while still maintaining the high degree of processing capability and the total flexibility achieved by using both types of  
25 processors on the same data.  
30

Texas Instruments TMS 320 DSP processors are disclosed in coassigned U.S. Patents 4,577,282 and 4,713,748 and in coassigned U.S. Patent application SN 025,417 filed March 13, 1987. Further background is disclosed in the  
35 publications Second Generation TMS 320 Users Guide and

Third Generation TMS 320 Users Guide from Texas Instruments Incorporated. These patents, said application and publications are hereby incorporated herein by reference.

## 5 Memory Structure

FIGURE 17 shows a view of the image processing system, as discussed with respect to FIGURES 1 and 2, showing a particular layout of memory. It should be kept in mind, however, that the particular memory sizes have been  
 10 selected for a particular project, and any type of arrangement of memory and memory capacities can be utilized with this invention. The parameter section of memory 10 can be incorporated within memory 10 or can be, if desired, a stand-alone memory. Under some conditions the parameter  
 15 memory need not be present depending upon the communication requirements of the processors.

## Crossbar Switch

FIGURE 18 shows the prioritization circuitry of  
 20 crossbar switch 20. Each vertical of the crossbar switch is connected in a round robin fashion to a prioritization circuit internal to the particular crosspoint. In every vertical the lowest horizontal, which is associated with the transfer processor, is not included in the  
 25 prioritization wiring. This is so that when none of the other horizontals in the same vertical have been selected, the transfer processor has access to the memory. The exact manner in which the prioritization circuitry operates and the manner in which the lowest horizontal operates will be  
 30 detailed more fully hereinafter with respect to FIGURES 19 and 20.

FIGURE 18 also shows the special situation of the instruction vertical I for the parallel processors. The instruction vertical for parallel processor 103 is  
 35 connected through crosspoint 4-7, which crosspoint is

09875436-060604  
 T09090-9875436

enabled by a signal on the SIMD lead via inverter 1801. This same signal is provided to every horizontal crosspoint 4-1 through 4-6 in the same vertical to render those crosspoints inactive. This signal and the manner in which the instruction vertical is connected to memory will be discussed hereinafter.

Turning now to FIGURE 19, the details of an exemplary crosspoint 1-5 is shown in detail. In the figure, the five sided box with a control line entering the side is a control switch, typically a FET device.

The functionality of the crosspoint logic is described. The crosspoint logic contains four functional blocks. These will each be described. The first functional block is address recognition block 1901 which compares five bits of the address supplied by the processor on bus 1932 with the unique five bit value of the memory module 10-15 (connected to crosspoint 1-5 via vertical 1 as shown in FIGURE 4) presented on bus 1930. The value presented on bus 1930 indicates the location of the memory within the address space. The comparison is achieved by five two-input exclusive-NOR gates 1920-1924 which perform individual bit comparisons. The outputs of these five gates are supplied to five of the inputs of the six input NAND gate 1910. The sixth input of gate 1910 is connected to the global access signal 1933 which indicates that a memory request is actually being performed and the address output by the processor should actually be compared. Only when signal 1933 is a logic one and the outputs of gates 1920-1924 are also all one will the output of gate 1910 be a logical zero. A logic zero indicates that a valid request for memory 10-15 is being made.

Digressing, a modification that can be made to this address recognition logic is to include a seventh input to gate 1910 (enable SIMD) that can be used as an enable signal for the crosspoint logic. A logical zero on the

5

10

15

20

25

30

35

The logic of the token latch records which crosspoint logic associated with memory 10-15 last gained access to the memory. This is indicated by a logical zero B1 signal being output by that crosspoint latch. The token latch logic works in conjunction with the prioritization block, to be described next, to cause the crosspoint which last accessed the memory to have the lowest priority access, if future multiple simultaneous accesses are attempted to the memory. How the token latch contents are altered will be described after the prioritization block has been described.

The prioritization block 1902 contains two two-input NAND gates 1911 and 1912. The two inputs of gate 1912 are supplied from the output of gates 1910 and 1911. The output of gate 1912 is signal A1 which connects to signal A of the vertically below crosspoint (1-4). One output of gate 1911 is the previously mentioned signal B which is connected to signal B1 from the token latch in the logic circuit associated with the next higher vertical (crosspoint 1-6). The other signal is also the previously described signal A which is connected to signal A1 from the prioritization block in the next higher vertical (crosspoint logic 1-6).

The prioritization logic forms a circular ripple path that begins with the crosspoint logic vertically below the last crosspoint to access the memory. This is indicated by a logical zero on a B1 signal. This causes the output of gate 1911 of the next vertical crosspoint below to be a logical one. This is gated by gate 1912 with the output of gate 1910 in order to produce signal A1. If the output of gate 1910 is a logical one, indicating that an address match by the address recognition logic wasn't found, then signal A1 will be a zero. This is passed to the next lower vertical crosspoint, causing its gate 1911 to output a one, and so on around the circular ripple path. If however the

T09090" 96T54860

output of gate 1910 is a zero, then the signal A1 will be output to the next crosspoint as a logical one. This, in conjunction with a one on all subsequent B inputs (since only the ripple start point can output a zero B signal),  
 5 causes all other gates 1911 around the ripple path to output logical zeros. Thus, a crosspoint can gain access to a memory only when it has a one on the output of its gate 1911 and it is producing a logical zero on the output of its gate 1910. This occurs only when an address match  
 10 is found by the address recognition block and the crosspoint is the first to request a memory access from the start of the circular ripple path.

The management of the token latch contents will now be explained. Gates 1913 and 1914 are designed to make sure  
 15 that the last crosspoint to gain memory access holds a zero in the token latch. Consider the following cases:

1. The token in token latch 1904 is a zero and no bus requires memory access. The zero ripples  
 20 completely around the circular carry path and returns to signal A of the originating crosspoint as a zero, causing the output of gate 1911 to be a one. The zero already held in the token latch (signal B1) causes the output of gate 1913 to be  
 25 a one. These two signals cause the output of gate 1914 to be a zero, which is loaded into the latch 1945/1946 by clock4, thus maintaining a zero in the token latch, thereby continuing the ripple propagation.  
 30
2. The token in token latch 1904 is a zero and one of the other crosspoints requires access to the memory. In this case, signal A will be received  
 35 back as a one, which in conjunction with the one on input B will cause the output of gate 1911 to

be a zero, causing the output of gate 1914 to be a one. This is then loaded into token latch 1904 by clock4 as a one. The token latch has thus become a one since another crosspoint has just gained memory access.

5

3. The token in token latch 1904 is a one and a crosspoint prioritized higher is requesting memory access. In this case A and B are both received as ones and, as in the above case, the token will similarly be loaded with a one.

10

4. The token in token latch 1904 is a one, the crosspoint is requesting memory access, and no higher priority crosspoint is requesting memory access. In this case either A or B will be received as a zero, causing the output of gate 1911 to be a one. The output of gate 1910 will be a zero, since the address recognition logic is detecting an address match. This will cause the output of gate 1913 to be a one. Since both inputs of gate 1914 are one, it will output a zero, which is loaded into token latch 1904 by clock4. The token latch has thus become a zero because it has just been granted memory access.

15

20

25

The fourth block of logic is the grant latch. The output of gate 1910 is passed through an inverter 1940 into one input of a two-input NAND gate 1915, whose other input is connected to the output of gate 1911. The one condition of a logical one on the output of gate 1911 and a zero on the output of gate 1910 causes the output of gate 1915 to be a zero. (Otherwise it is a one). This condition occurs when the crosspoint is successfully granted access to the memory, and can occur on only one of the crosspoints

30

35

 05875136-060604  
 T09099" 9ET 54850



associated with the memory. The output of gate 1915 is loaded into latch 1941/1942 through transistor 1951 by clock1. (In practice clock1 and clock4 will operate together so that the token latch and the grant latch are updated together). The output of gate 1942 is loaded through transistor 1952 by clock2 into latch 1943/1944. The output of gate 1944 is passed to gate 1949 which produces the connect signal to the crosspoint switches 1905, which connect processor bus 1932 to memory bus 1931. These crosspoint switches can be individual n-type transistors in their simplest implementation.

The output of gate 1942 is also supplied to the gate of transistor 1958 which connects between signal 1934 and the source of transistor 1957, whose drain connects to ground, and whose gate is connected to clock2. Transistors 1957 and 1958 cause signal 1934 to be connected to ground when the crosspoint has successfully been granted memory access. This indicates to the processor that it can proceed with the memory access. If however signal 1934 does not go low when a memory access is attempted, then another crosspoint has gained memory access and the processor must halt and re-request access to the memory. The round-robin prioritization scheme described ensures that only a limited number of retries need be performed before access is granted.

An example of the timing of the crossbar signals is given in FIGURE 20. In this figure PP2 and PP3 are both trying to access the same RAM every cycle, but the round-robin priority logic causes them to alternate. PP2 is calculating and outputting addresses S, T and U, and PP3 is calculating and outputting addresses V and W. It can be seen from the :5 MS address" signals how the GRANTED-signal is used to multiplex between the last address (in the case of a retry) and the new address being calculated. The PPs assume that if the GRANTED- signal is not active by

the end of the slave phase then contention occurred, and the master update phases of the fetch, address and execute pipeline stages are killed.

## 5 Integration of the Switch Matrix

As discussed herein, memory contention is handled by a token passing arrangement having logic circuitry individual to each crosspoint. In one embodiment, the logic circuitry is positioned in direct association with each crosspoint. Thus, since the crosspoints are spatially distributed across the substrate in conjunction with their respective ports, the contention control logic is likewise distributed spatially. In addition to saving space the actual logic of the circuit can grow as the switch grows. In this manner the logic can be positioned in one of the layers of the silicon so that no additional silicon chip area is consumed. This has the advantage of conserving space while also minimizing connections to and from the token passing circuit.

## 20 Synchronized MIMD

Each processor 100-103, as shown in FIGURE 21, has associated with it a register 2100-2103 respectively for indicating if synchronized operation is required. Also included, as will be seen, is a register for holding the address (identity) of the other processors synchronized with that processor. The instruction stream contains instructions which indicate the beginning and end of a series of instructions that must be executed in synchronization with the processors. Once the code for starting a synchronized instruction stream arrives at a processor, that processor, and all the processors in the synchronized set, can only execute instructions in lock step with each other until such time as the end of synchronized code instruction is encountered.

Using this approach, no messages need be transferred between processors, and the processors will remain in step for one cycle, or a number of cycles, depending upon the instruction stream being executed. No external control, other than the instruction stream, is required to establish the synchronization relationships between processors.

Turning to FIGURE 22, within each parallel processor 100-103, there is a sync register 2207 containing four bits labelled 3, 2, 1, 0 that relate to processors 103, 102, 101 and 100 respectively. One bit relates to each processor 100-103. The other processor(s) to which a particular processor will synchronize is indicated by writing a one to the bits corresponding to those processors. The other processor(s) which are expecting to be synchronized will similarly have set the appropriate bits in their sync register(s).

Code that is desired to be executed in synchronization is indicated by bounding it with LCK (Lock) and ULCK (Unlock) instructions. The instructions following the LCK, and those up to and including the ULCK, will be fetched in lock-step with the other parallel processor(s). (There must, therefore, be the same number of instructions between the LCK and ULCK instructions in each synchronized parallel processor).

It is more usually synchronized data transfer that is required rather than synchronized fetching of instructions. It is a consequence of the parallel processors' pipelines however that the transfer(s) coded in parallel with the LCK instruction and those up to and including the instruction immediately preceding the ULCK instruction, will be synchronous. They may not necessarily (due to memory access conflicts) occur in exactly the same machine cycle, but the transfers coded in the following instruction will not proceed until all the synchronized transfers of the previous instruction have occurred. The order of the load

The knowledge that synchronized code is being executed is recorded by the S (synchronized) bit in each status register. (This bit is not actually set or reset until the master phase of the address pipeline stage of the LCK or ULCK instructions, respectively, but the effect of the LCK or ULCK instruction affects the fetch of the next instruction during the slave phase). This bit is cleared by reset and by interrupts once the status register has been pushed.

Turning now to processor 100, it should be noted that the presence of a 0 in any bit of sync register 2207 causes a logic one to appear on the output of the respective NAND gate. Thus, with the example shown, the NAND gates 2203 and 2204 would have logic ones on their respective output. These ones are supplied to the input of NAND gate 2206. NAND gate 2206 will not allow processor 100 to execute any more instructions of code until all of its respective inputs are one. Note that the presence of the zeros in the bit positions 1 and 2 of register 2207 causes the respective gates 2203 and 2204 to ignore the presence of any signals on leads 1 and 2 of bus 40. Thus, the execution of code was controlled by gate 2206, in this case in response to the information on leads 0 and 3 of bus 40.

The lock instruction will cause the S bit to become set  
 which is a logic 1 to one of the inputs to gate 2201. For  
 the moment we will ignore the presence of the okay to sync  
 signal which is a signal which controls the timing of the  
 5 actual execute for the processor. The output of gate 2201  
 for each of the processors sync registers is connected to  
 a different lead. Thus, gate 2201 from processor 100 is  
 connected to lead 0, while gate 2201 from processor 101 is  
 connected to lead 1, etc. Note that the output of gate  
 10 2201 from processor 100 is connected to the 0 input of  
 gates 2205 of all of the other processor registers. Since  
 in processor 101 and 102, gates 2205 are connected to logic  
 zero, this has no effect. However, in processor 103 where  
 15 gate 2205 is connected to a logic 1 of the register, it is  
 thus controlled by the output on lead 0 of bus 40 which in  
 fact is controlled by the output of gate 2201. Thus,  
 processor 103 is controlled by the actions which occur  
 within processor 100, which is exactly what we desire if  
 processor 103 is to be synchronized with processor 100. A  
 20 review of the circuitry would show that the same function  
 operates in reverse from processor 103 to processor 100  
 since in processor 103 gate 2201 is associated with lead 3  
 of bus 40, which in turn is associated with gate 2202 of  
 processor 100, which in turn is also controlled by a one in  
 25 sync register 2207.

Now returning to the signal on gate 2201 which is the  
 okay to sync signal. When that signal goes to logic 1,  
 then it is okay to execute code, and all of the other  
 processors having a one in the sync register bit 0 position  
 30 of the respective register will operate in synchronization  
 with that signal. Thus, if the okay to sync signal goes  
 low signifying a problem with the cache memory or any other  
 problem with the execution of code, all of the processors  
 synchronized therewith will wait until the problem is  
 35 clear. Thus, we have full synchronization between

TOP SECRET

processors as controlled by the codes periodically stored in the respective registers. All of the processors can be synchronized or any combination of processors can be synchronized with each other, and there can be any number  
 5 of different synchronizations occurring between processors.

Since it is the instruction fetch that is synchronized, it is possible to interrupt synchronized code. This will immediately cause the parallel processor's sync signal to become inactive. Cache misses and  
 10 contention will have a similar effect, keeping the machines in step. In the case of contention, however, the two instructions following the one experiencing contention will have already been fetched into the pipeline before the pipeline pauses.

15 It is possible to put idle instructions into synchronized code, thus pausing the operation of all the synchronized parallel processors until a particular parallel processor has been interrupted and returned from its interrupt routine.

20 Since it is necessary to be able to interrupt synchronized code, any instruction that specified the program counter PC in any one processor as a destination will immediately disable the effect of the S bit of the status register (with the same timing as the ULCK  
 25 instruction), but the S bit will remain set. Once the two delay slot instructions have completed, the effect of the S bit is re-enabled. This mechanism prevents problems with being unable to interrupt synchronized delay slot instructions. The sync logic therefore treats branches,  
 30 calls and returns (implemented as a PC load followed by two delay slot instructions) as a single instruction. The sync signal will be driven inactive during the two delay slot instructions and they will be fetched without looking at the sync signals. If a LCK instruction is put in a delay  
 35 slot, it will take effect after the delay slot instructions

TOP SECRET

An example of how synchronization works is given in  
5 FIGURE 23. In this case, parallel processor 2 and parallel  
1 exchange the contents of their data D0 registers (FIGURE  
33), assuming that A0 and A1 contain the same addresses in  
each parallel processor. It also assumes that A0 and A1  
point to different RAMs to avoid contention. (It would  
10 still work if they pointed to the same RAM, but would take  
extra cycles).

Synchronization in SIMD is implicit, so the LCK and  
25 ULCK instructions have no purpose and so will have no  
effect if coded. The S bit in the status register will  
have no effect if anyone should set it to one.

The instructions shown in the appendix (LCK) is used to begin a piece of MIMD synchronized parallel processor code. It will cause the parallel processor to wait until all the parallel processors indicated by ones in the sync register are in sync with each other. The following instructions will then be fetched in step with the other MIMD parallel processors. Execution of the address and execute pipeline stages will occur as each successive

instruction is synchronously fetched. The S bit of the status register is set during the address pipeline stage of this instruction.

The instruction shown in the appendix (ULCK) unlocks the MIMD parallel processors from each other. They then resume independent instruction execution on the next instruction fetch.

### Sliced Addressing

10 Sliced addressing is a technique for taking adjacent information from one memory space and distributing it in a manner to a number of separate different memory spaces so that the information when it has been distributed can be accessed simultaneously by a number of processors without  
15 contention.

As an example, reference is made to FIGURE 24 where there is shown an external image memory buffer 15 having a row of adjacent pixels numbered 0-127, and this row has the letter "a" referencing it. This information is  
20 transferred, using the sliced addressing technique, via bus 2401, into memory subsystem 10 whereby the first sixteen pixels (0-15) are placed into the first memory 10-0 referred to by address 0-15. Then the next sixteen pixels are placed into memory 10-1. In this example this process  
25 is continued through eight memories such that pixels 112-127 are placed into final memory 10-7. The sliced addressing logic 2401 is implemented in the transfer processor and also in the crossbar address units of the parallel processors which will be described hereinafter.

30 The prior art means of address calculation would produce in the given example 128 consecutive addresses. This would mean that the data would be placed within one memory. In the given example the data would appear at consecutive addresses within memory 10-0. This would not  
35 allow a number of processors simultaneous access to that



information without contention since they would all be trying to access the same memory. Thus, in the prior art, pixels 0-15 would be in row A of memory 0 with bits 16-31 in row B and bits 32-47 in row C, etc., until all of the 127 adjacent pixels would be in various rows of memory 0. Since the various different processors are working in parallel to process information, they could all contend for access to memory 0 to various pixel bytes, and accordingly time would be wasted, and the value of the parallel processing would be mitigated.

FIGURE 25 shows a prior art adder which is used for controlling the location of the address for various bits. FIGURE 25 shows three single bit adders 2501, 2502, 2503, which are part of a full adder having a number of single bits equal to the address range of the memory. These adders work such that one bit of the address is provided to each A input of the various adders 2501-2503. The least significant bit of the address would go to adder 2501, and the most significant bit would go to the highest single bit adder 2503.

The B input receives the binary representation of the amount to be indexed for the address for storage purposes. The combination of adders 2501-2503 will produce a resulting address which is used for accessing memory. Each individual adder will output a carry signal to the next highest numbered adder carry input signal. Each individual adder bit will take in the three inputs A, B and carry in, and if there are two or three ones present on any of those inputs, then the carry out from that cell will be a one. This is supplied to the next most significant carry in input of the adder. This process is repeated for each individual adder bit to produce a resultant address of the size required to access the memory space. The fact that each carry out connects directly to the next most significant carry in, means that the resultant address is

always part of a contiguous address space. In the previous example, if an index of value one is supplied to the B inputs of the adder, then the resultant address output to memory will be one greater than the original address supplied on the A inputs.

With reference to FIGURE 26, the modification to the previously described normal adder is made whereby the carry out of each cell is multiplexed with the carry in signal supplied to each cell, such that the signal that is passed to the next most significant carry in inputs of the adder can be selected to be either the carry out of the previous cell or the carry in for that previous cell. As an example, consider cell 2505. Its carry out signal is supplied to the multiplexer 2508, and the multiplexer's other input is the carry in signal to 2505. Signal B is used to control the multiplexer causing either the carry out or the carry in of cell 2505 to be passed on the carry in input of the next most significant cell.

Another modification to the standard cell is to include a control input labelled  $\overline{ADD}$  which is supplied by the same control signal that controls the multiplexer signal B. If a logical one is supplied on signal B, then the carry in signal of 2505 is supplied to the carry in signal of the next most significant cell. The presence of a logical one on signal B also inhibits the add function of cell 2505 such that the original address supplied on input A is passed straight through to the output without modification. This has the effect of protecting the address bit associated with the presence of a one on input B. It can be seen that by supplying a number of ones to the control signals of the modified adder, the carry out of a cell from the least significant bit can be propagated a number of cells along the length of the adder before being supplied to the carry in of a cell which will perform the add function. This would be the next most significant cell

which had a zero on the  $\overline{ADD}$  control signal. The effect of this is to protect the address contained within the cells which have been bypassed so that a number of bits of the address range have been protected from modification. With  
 5 reference to the previously described example, by supplying ones on the multiplexer and  $\overline{ADD}$  control signals, an address increment from pixel 15 in memory 0 can be made to pixel 16 in memory 1 so that the memory can be addressed as one continuous address space. The multiplexer control signals  
 10 are referred to as a sliced mask because they will mask out certain bits from the address range and cause the data which has been distributed in memory to be accessed as a slice indicated in FIGURE 24.

It should be noted that this circuitry is used both  
 15 for storing adjacent information or for retrieving adjacent information. Also, some information should be provided and stored in the same memory and should not be sliced, and this is denoted by providing all zeros to the ABC leads of the slice mask. When this occurs, the individual adders  
 20 2504-2506 act in the same manner as the prior art adders 2501-2503. It is also important to keep in mind that there are different types of distributed data that should be sliced across several memories and not just pixel information. This would occur anytime when it is  
 25 conceivable that several processors would be accessing the same type of information at the same time for whatever processing would be occurring at that point.

It is also important to keep in mind that to distribute memory as disclosed in the sliced addressing  
 30 mode does not in any way waste memory because the rows B and C which are not used for the particular pixel or other information to be stored would be used for other information. The only "penalty" that conceivably could occur is the additional chip space required to construct  
 35 the multiplexers and the additional interconnections of the

For example, if information which at one time is sliced because it is being used in a parallel processing mode is later determined to be used for a single processor for a single period of time, it would be advantageous to provide all zeros on the slice mask for that time period thereby storing the information in a single memory so that a single processor can then access the single memory, in this way again gaining valuable time over the slice method. This then gives a high degree of flexibility to the design of the system and to the operational mode for storing data.

Turning now to FIGURE 27, an example of the way in which a typical quantity of pixels may be distributed over a number of memories is shown. In this example each individual memory is two kilobytes in size, and the start and end addresses of each of these memories are indicated. For example, memory 0 begins at address all zeroes and finishes as address 07FF. Memory 1 begins at 0800 and ends at 0FFF and so on through to memory 7 which begins at 3800 and ends at 3FFF. A quantity of pixels are shown distributed in a slice across these memories, 64 pixels per memory. Consider for a moment stepping through the 64 pixels within the slice of memory 3. We can see that the pixels are arranged from addresses 1900-193F. The next adjacent piece of information is not resident at the next

30

Before beginning a detailed description of how the MIMD/SIMD operational modes change the reconfigure of the memory, it would be good to review FIGURE 4 with respect to the processors' memory and crossbar interconnections

35

5

10

20

35

5

10

15

20

25

30

35

processors for communicating the external address of the instructions required to be executed, and also the place within the instruction memory 10-5, 10-9 or 10-13, respectively, where the next sequence of instructions are to be stored. Once the transfer processor has fetched these instructions, an acknowledged signal is passed to the parallel processors from the transfer processor indicating that the code has now been fetched. The parallel processor can then perform instruction execution, again from the memory until such occasion as the instruction stream is found to be absent and the process is again repeated.

In the SIMD configuration because processors 101, 102 and 103 are executing from the same instruction stream, the cache logics 2801, 2802 and 2803 within the processors are disabled because they perform no function. The program counters 2811, 2812 and 2813 contents are irrelevant because they perform no purpose in fetching instructions because in the SIMD configuration all instructions are fetched by processor 100. In the SIMD configuration, therefore, it is desirable to use memories 10-5, 10-9 and 10-13 for storing data. In order to do this, crosspoints 14-1 through 14-6, 9-1 through 9-6 and 4-1 through 4-6 are enabled, thus allowing those memories to be accessed by the processors for data. This means that the memory utilization in the system is maintained at its optimum level for both SIMD and MIMD configurations.

#### Imaging Personal Computer

The imaging personal computer (PC) shown in FIGURES 46-52, can be constructed of three major elements, a camera sensing device 4600, shown in FIGURE 46, an imaging processing device 4602 and a display device 4801 (FIGURE 48). The imaging PC is not restricted to the use of a camera 4600 or a display 4801 and many forms of image input/output can be used.



Camera 4600 could be focused in front of screen 4601 of the PC and a hand 4603 can be used to input information by "signing" as typically done for deaf communication. The "signing" could be observed by the camera, and the screen could be used to display either the sign "two" or can be used to further process the information as discussed previously with respect to FIGURE 11. The output bus from the PC could also contain the digital representation of the information being input via camera 4600, in this case the binary bits representing two. Thus, the user could utilize spreadsheets and other information obtaining information both from a keyboard or other traditional manner in ascii code as well as from a visual or video source such as camera 4600 or video recorder device or any other type of video input using an imaging code input. The video input can be recorded on tape, on disc or on any other media and stored in the same manner as information is currently stored for presentation to a PC.

Some of the features that an imaging PC can have are

- 1) acquiring images from cameras, scanners and other sensors;
- 2) understanding the information or objects, in a document;
- 3) extracting pertinent information from a document or picture;
- 4) navigating through a data base combining images as well as textual documents;
- 5) providing advanced imaging interfaces, such as gesture recognition.

The PC can be used to create instant data bases since the information put into the system can be read and the informational content abstracted immediately without further processing by other systems. This creates a data base that can be accessed simply by a match of particular words, none of which had been identified prior to the storage. This can be extended beyond words to geometric shapes, pictures and can be useful in many applications. For example, a system could be designed to scan a catalog, or a newspaper, to find a particular object, such as all of





FIGURE 50 describes an imbedded application of the image system processor 5000. In this case images are sensed again via CCD's 5004 or other sensors which collect information from the world, such as the presence of an intruder in a security application. This information is placed in a frame buffer or VRAM 5003 which is the external memory for the image system processor 5000. Alternatively, the ISP can be used as a pattern (or person) recognizer and output control information fed to latch 5008. This information would be used to control a mechanism 5009, such as a door lock or factory process or the like. Also, the output from latch 5008 could be presented to output display 5010. The program or instructions have been previously stored in a hard drive or optical disc 5002 or 5001. These devices can also be used to store incidences of information such as again in a security application, the image of an intruder. The statistical accumulated record keeping 5007 maintains system status or occurrence of events which have occurred.

FIGURE 51 depicts a handheld imaging PC. In this case the image system processor 5106 accepts input from two charge couple devices 5105 which provide position input which is then processed to extract user supplied gestures and control of the PC. The position and orientation of the user's hand or pseudo pen would then be used to control the device or in conjunction with the ISP to extract meaningful messages or characters. Flat panel display 5104 provides an output information display of this handheld PC. Optionally, an external camera 5103 would allow the user to collect images outside of the scope of the handheld PC's memory. A host or printer port would also be provided to allow the user to download or print information contained in the handheld PC.

FIGURE 52 describes an application of the ISP in a network configuration with a host 5205 which provides

necessary image information collected off-line either remotely or in some central office and then distributed to buffer 5201 which is then used by the imaging PC configuration to provide information to the image system processor 5200. An alternative method of obtaining information is via scanner 5207 working in conjunction with front end processor 5206. This reduced cost version of the imaging PC (with respect to FIGURE 49) would permit the resource sharing by networking image collection devices. A printer port would also be provided via printer interface 5203 and its connection to printer mechanism 5204 which would allow the user to print the compound documents which contain the normal textual and graphic information in addition to images or enhanced images via the image system processor.

The compact structure of the image processing system, where all of the parallel processing and memory interaction is available on a single chip coupled with a wide flexibility of processor memory configurations and operational modes, all chip controlled, contributes to the ability of the imaging PC to accept image data input as well as ascii input and to allow the two types of data to be simultaneously utilized.

#### 25 Ones Counting Circuit

FIGURE 53 shows an imaging system 5310 operable to process image data using combinations of various processing algorithms. An imaging device 5312, such as a video camera, a still image camera, a bar code reader and the like, is used to capture images and provides them to an image data memory 5314. The captured images are stored in image data memory 5314 until they are accessed by an image processor 5316 addressed by an address generator 5318. Image processor 5316, such as the processor shown in FIGURES 1 and 2, performs signal processing functions

including statistical processes on the image data, such as histograms. A ones counting circuit 5320 is provided to generate a count of the number of "ones" in the image data. Information, such as the number of "ones" along a  
 5 projection line in the image data, is used to provide a statistical analysis of the image data, which may be used for pattern recognition. The histogram of the image data may be compared to predetermined image patterns to recognize a pattern match. An output device 5322 is  
 10 coupled to image processor 5316 and is available for displaying any output of imaging system 5310. The output device 5322 may be a monitor or a hard copy generating device.

It should be understood that the overview of the  
 15 imaging system 5310 described above provides an example of an environment in which the present invention may advantageously operate, and the description above in no way limits the applicability of ones counting circuit 5320.

Referring to FIGURE 54, a logic gate level  
 20 implementation of a ones counting circuit 5320 is shown. The ones counting circuit 5320 consists of a matrix 5424 having M number of rows and N number of columns of count cells 5426a through 5426l, where M is equal to three and N is equal to four in the embodiment shown in FIGURE 54. For  
 25 an input binary string of  $X_n$  number of bits, M may be determined by:

$$M = \log_2 (X_n + 1)$$

30 rounded up to the nearest integer, and N may be determined by:

$$N = X_n.$$

The matrix 5424 receives a binary string denoted by X, and  
 35 produces a binary number denoted by Y, indicative of the

number of "ones" in the binary string. Another output, denoted by Z, is used in a minimized ones counting circuit matrix, to be discussed in detail below.

15       An XOR gate generates a logic level "one" at an output only if an odd number of "ones" are present at its input. For example, XOR gate 5430a will produce a "one" at output 5438a if a "one" is present at only one of its inputs 5440a and 5442a.

Count cells 5426a through 5426l are arranged in rows and columns in matrix 5424. The interconnections of count cells 5426a, 5426b and 5426e will be used to illustrate the interconnections of the whole matrix 5424. As shown in FIGURE 54, count cell 5426a is arranged to be left of count cell 5426b and above count cell 5426e. Count cell 5426a is connected to count cell 5426b, where output 5438 of XOR gate 5430 of count cell 5426b is connected to inputs 5432a and 5440a of count cell 5426a. Count cell 5426a is

connected to count cell 5426e, where output 5436 of AND gate 5428 of count cell 5426e is connected to both input 5434a of AND gate 5428a and also input 5442a of XOR 5430a gate of count cell 5426a. The interconnections just  
 5 described may be expanded to the whole matrix by using the connection between count cells 5426a and 5426e for inter-row connections and the connection between count cells 5426a and 5426b for inter-column connections. It is appropriate to note that matrix 5424 may be implemented  
 10 with the rows as the columns and vice versa, and the matrix itself transposed without departing from the teachings of the present invention.

In order to further describe the structure of matrix 5424, the following convention will be used when referring  
 15 to the rows and columns: rows have row numbers zero through  $(M - 1)$ , and columns have column numbers zero through  $(N - 1)$ , where the bottommost row is row zero and the right most column is column zero. In the example shown in FIGURE 54, M is three and N is four. Furthermore, references may  
 20 be made to a count cell at a position  $(x,y)$ . The coordinates x and y indicate the column number and row number, respectively, of a count cell. For example, count cell 5426a is at position  $(3,2)$ .

Accordingly, matrix 5424 comprises interconnected  
 25 count cells 5426a through 5426l arranged in rows and columns where row zero receives the binary string X, row one receives the AND gate outputs of row zero, and row two produces output Z. Column-wise, column zero receives "zeros" from any source to begin the propagation, column  
 30 one receives the XOR gate output of row zero, column two receives the XOR gate output of row one, and column three produces output Y indicative of the number of "ones" in binary string X. The logic level "zero" received by column zero may be produced by hardwiring the inputs to ground.

109090" 95752880



For the purpose of illustration, a binary string 1101 ( $X_3 = 1$ ,  $X_2 = 1$ ,  $X_1 = 0$ ,  $X_0 = 1$ ) is received by row zero of matrix 5424. AND gate 5428 of count cell 5426l produces a "zero" at its output, and XOR gate 5430 of count cell 5426l produces a "one" at its output. The logic level "one" from XOR gate 5430 of count cell 5426l is propagated down row zero, and the outputs of the XOR gates of each cell toggles each time there is a "one" in the corresponding X input. Therefore, the output of XOR gate 5430 of count cell 5426k remains at logic level "one," the output of XOR gate 5430 of count cell 5426j toggles to a "zero," and the output of XOR gate 5430 of count cell 5426i toggles again to a "one." This produces a "one" at the output of row zero, which makes  $Y_0$  equal to "one."

In row one, the XOR gates toggle their outputs in a similar fashion. The output of XOR gate 5430 count cell 5426h is a "zero," having received a "zero" from AND gate 5428 of count cell 5426l. The output of XOR gate 5430 of count cell 5426g remains at logic level "zero," having received "zeros" from both XOR gate 5430 of count cell 5426h and AND gate 5428 of count cell 5426k. Subsequently, the output of XOR gate 5430 of count cell 5426f toggles to a "one," having received a "zero" from XOR gate 5430 of count cell 5426g and a "one" from AND gate 5428 of count cell 5426j. The output of XOR gate 5430 of count cell 5426e again toggles, to a "zero" this time, having received a "one" from XOR gate 5430 of count cell 5426f and a "zero" from AND gate 5430 of count cell 5426i. As a result, a "one" is produced at the output of row one, which makes  $Y_1$  equal to "one."

In row two, the output of XOR gate 5430 of count cell 5426d is a "zero," having received the hardwired zero and another "zero" from AND gate 5428 of count cell 5426h. The output of XOR gate 5430 of count cell 5426c remains at logic level "zero," having received "zeros" from both XOR

gate 5430 of count cell 5426d and AND gate 5428 of count cell 5426g. Subsequently, the outputs of XOR gates 5430 of both count cells 5426a and 5426b also produce "zeros," which produce a "zero" at the output of row two, making  $Y_2$  equal to "zero." Therefore, for the example input binary string  $X = 1101$ , the output is the binary number  $Y = 011$ , which is three. Indeed, there are exactly three "ones" in the example binary string input  $X = 1101$ .

It can be appreciated that the ones counting circuit 5320 is an asynchronous circuit, which receives inputs and generates outputs without requiring clock signals. Thus, in matrix 5424, an output is available as soon as the inputs are received and the signals are propagated through the matrix. The longest propagation time through the matrix would be the time it takes for the signals to propagate through the longest path which includes count cells 5426l, 5426h, 5426d, 5426c, 5426b, and 5426a.

Matrix 5424 shown in FIGURE 54 is rectangular and comprises identical count cells 5426. These characteristics make the ones counting circuit compact and easily laid out for semiconductor mask production. However, matrix 5424 may be minimized by using fewer count cells and/or fewer gates.

Referring to FIGURE 55, a minimized ones counting circuit matrix 5544 for a four bit binary string input is shown. Matrix 5544 includes interconnected count cells 5546a through 5546e. For a minimized matrix, the number of rows  $M$ , and the number of count cells in each row  $N$ , are determined as follows:

30

$$M = \log_2 X_n$$

rounded up to the nearest integer, and for each row

35

$$N = X_n - 2^r,$$

where  $X_n$  is the number of bits in the input binary string  $X$ , and  $r$  is the row number ranging from zero to  $(M - 1)$ . In the example shown in FIGURE 55, the number of bits  $X_n$  of the  
 5 input binary string  $X$  is four. Using the above formulas, the number of rows,  $M$ , is equal to two. To calculate for  $N$  for the first row,  $r$  is equal to zero, which makes  $N$  equal to three. For the second row,  $r$  is equal to one, which makes  $N$  equal to two. Thus, a minimized matrix of  
 10 three count cells in the first row and two count cells in the second row, totaling five count cells, can compute the number of "ones" in a four bit binary string, as compared with the twelve count cells in the full matrix 5424 (FIGURE 54).

15 Each count cell 5546a through 5546e comprises an AND gate 5548 coupled to an XOR gate 5550, identical to the count cells of the full matrix 5424 shown in FIGURE 54. The binary input string  $X$  is received by the inputs to count cells 5546c through 5546e; the output binary number  
 20  $Y$  is produced at the outputs of count cells 5546a and 5546c.

In the example shown in FIGURE 55,  $X_3$  is received by the inputs 5552 to AND gate 5548 and XOR gate 5550 of count cell 5546c;  $X_2$  is received by the inputs 5554 to AND gate  
 25 5548 and XOR gate 5550 of count cell 5546d.  $X_1$  is received by the inputs 5556 to AND gate 5548 and XOR gate 5550 of count cell 5546e;  $X_0$  is received by the other inputs 5558 to AND gate 5548 and XOR gate 5550 of count cell 5546e.

The most significant bit of the binary number output  
 30  $Y$ ,  $Y_2$ , is produced at output 5560 of AND gate 5548 of count cell 5546a.  $Y_1$  is produced at output 5562 of XOR gate 5550 of count cell 5546a. The least significant bit  $Y_0$  is produced at output 5564 of XOR gate 5550 of count cell 5546c.

Because the minimized matrix 5544 is not rectangular, the interconnections between the count cells are modified. In particular, if a count cell at position  $(x,y)$  is not present as compared with the full matrix, the count cell in the row immediately "below" it is connected to the input of the XOR gate of the count cell  $(x+1,y)$  immediately to the left of the missing cell. If more than one count cell is absent, for example, count cells at positions  $(x,y)$  and  $(x+1,y)$ , then only the output of the AND gate of the count cell at position  $(x+1,y-1)$  need to be connected to the input of the XOR gate of the count cell at position  $(x+2,y)$ . In the embodiment shown in FIGURE 55, the count cells at positions  $(0,1)$  and  $(1,1)$  are absent, so the output of AND gate 5548 of the count cell 5546e at position  $(1,0)$  is connected to the inputs of AND gate 5548 and XOR gate 5550 of count cell 5546b at position  $(2,1)$ . Further, the count cell at position  $(0,0)$  is also absent as compared with the full matrix implementation. The input  $X_0$ , then, is directly connected to inputs 5556 and 5558 of AND gate 5548 and XOR gate 5550, respectively, of count cell 5546e at position  $(1,0)$ . The count cell at position  $(3,0)$  is also absent, so the output  $Y_2$  is directly provided by the output 5560 of AND gate 5548 of count cell 5546a at position  $(3,1)$ .

Using the prior example  $X = 1101$ , where  $X_3 = 1$ ,  $X_2 = 1$ ,  $X_1 = 0$ , and  $X_0 = 1$ , the output of AND gate 5548 of count cell 5546e is a "zero," and the output of XOR gate 5550 of the same count cell 5546e is a "one." The logic level "one" from XOR gate 5550 of count cell 5546a is propagated down row zero and the outputs of XOR gates of each cell toggle each time there is a "one" in the corresponding  $X$  input. Therefore, the output of XOR gate 5550 of count cell 5546d toggles to a "zero," and the output of XOR gate 5550 of count cell 5546c toggles again to a "one." This

produces a "one" at the output of row zero, which makes  $Y_0$  equal to "one."

In the second row, the output Z of AND gate 5548 of count cell 5546b is a "zero," having received a "zero" from  
 5 AND gate 5548 of count cell 5546e. XOR gate 5550 of count cell 5546b outputs a "one," having received a "zero" from count cell 5546e and a "one" from count cell 5546d. XOR gate 5550 of count cell 5546a outputs a "one," having received a "zero" from count cell 5546c and a "one" from  
 10 count cell 5546b. This produces a "one" at the output of row one, making  $Y_1$  equal to "one." In addition,  $Y_2$ , which is the output of AND gate 5548 of count cell 5546a, is a "zero." Therefore, the output binary number is equal to  $Y = 011$ , indicating that there are three "ones" in the  
 15 input binary string  $X = 1101$ .

Matrix 5544 may be further minimized by eliminating some logic gates, such as AND gate 5548 of count cell 5546b, shown in broken outline. Since the output Z of AND gate 5548 is not required to assemble output binary number  
 20 Y, AND gate 5548 can be eliminated. Therefore, in a minimized matrix, AND gates of count cells immediately adjacent to absent count cells in the same row may be removed to further reduce the size of the ones counting circuit.

25 It can be appreciated that the present invention is not limited in scope to the circuit implementation described and shown herein. In particular, alternative embodiments may include circuit implementations derivable from the present embodiment by Boolean logic as known in  
 30 the art. For example, an AND gate such as AND gate 5548 may be equally implemented by a NAND gate coupled to an inverter. Furthermore, by De Morgan's theorem as known in the art, an AND function may be implemented by an OR gate with an inverter coupled to its output and with the input  
 35 signals to the OR gate inverted. Such alternate circuits

derivable from the present embodiment are within the scope of the invention.

Referring now to FIGURE 56, an example application in character recognition of the present invention is shown.

5 A matrix of pixels 5666 consists of "zeros" and "ones" forming a letter "F." The pixels 5666 may be gathered by an aforementioned imaging device and stored in an image data memory. The matrix of pixels 5666 is processed row-wise and column-wise to generate row counts 5668 and column  
10 counts 5670 of the number of "ones" present in each row and column, respectively. The row counts 5668 are generated by providing each row of the pixel matrix 5666 as binary string input X to the ones counting circuit. Thus, a count of the number of "ones" of each row is generated. In the  
15 example shown in FIGURE 56, the capital letter "F" has no "one" pixels in the first two rows. In row three, there are four "ones" forming the first horizontal line in the letter. In row four, there is only one "one". Row five has three "ones" which form the second horizontal line in the letter "F." In each of rows six and seven, there is  
20 one "one".

Similarly, column counts 5670 are generated by providing each column of the pixel matrix 5666 to the input of the ones counting circuit. Columns one and two contain  
25 no "ones." In column three, there are five "ones" forming the vertical line in the letter "F." In column four, there are two; in column five, there are also two; in column six, there is one; and in columns seven and eight, there are none.

30 Therefore, the row counts and column counts of all characters and any image pattern may be generated and stored as histograms in a pattern recognition system, so that they may be used as a standard for comparison against new character image samples.

While the preferred embodiment of the present invention counts the number of "ones" in an input binary string, it is conceivable to implement a "zero" counting circuit operable to count the number of "zeros" in a binary string in an alternate embodiment by adding inverters at the input of the ones counting circuit matrix. Such a "zero" counting circuit is an alternate embodiment and is within the teachings of the present invention.

Although the present invention has been described in detail, it should be understood that various changes, substitutions and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.

## 15 PROCESSOR DETAILS

The following discussion pertains to the master processor, the parallel processors, and the transfer processor as detailed in FIGURES 29-45. While not necessary for an understanding of the operation of the invention claimed, this discussion may be helpful to give a specific embodiment of many such embodiments. The precise system used will depend upon the system requirements and can, in fact, vary substantially from the following discussion.

### 25 Parallel Processor

#### Master Processor

Turning now to FIGURE 29, we can look at workings of master processor 12 which serve to control the operation of the entire image system processor including controlling the synchronization and other information flowing between the various parallel processors. Master processor 12 executes instructions which can be 32 bit words having opcodes controlled by opcode circuit 2911 and register file 2901. Program counter 2903 operates under the control of

control logic 2904 to control the loading of instructions from bus 172 into opcode register 2911. Control logic 2904 then decodes the instruction and controls the operation on master processor 12 based on the information presented.

5        In addition to integer execution unit 2902, there is a floating point execution unit comprised of two parts. Part one is a floating point multiplier comprised of multiplier 2905, normalized circuit 2906 and exponent adder 2907. Part two is a floating point adder comprised of  
10        prenormalizer 2908 and 2909 and postnormalizing shifter 2910.

Program counter register 2903 is used to provide the address output along bus 172 when it is required to read 32 bit instructions. Acting in accordance with the  
15        instructions decoded from opcode register 2911, integer execution unit 2902 can provide addresses which are output over bus 171 to control the reading of data from a data cache external to the master processor. Data is returned over the data part of bus 171 and stored in register file  
20        2901.

The Instruction bus 172 and data bus 171 each consist of an address part and a data part. For instruction bus 172, the address part comes from the program counter 2903 and the data part is returned to opcode register 2911. For  
25        data bus 171, the address part comes from the output of the integer ALU 2902 and the data either comes from register file 2901 if it is a write cycle or is returned to register file 2901 if it is a read cycle.

The manner in which the various elements of master  
30        processor 12 interact with each other are well-known in the art. One example of the workings of a graphics processor is shown in copending U.S. patent application of Karl Guttag, David Gulley, and Jerry Van Aken, entitled "Graphics Processor Having a Floating Point Coprocessor",



Serial Number - 387,472, filed July 28, 1989, which application is hereby incorporated by reference herein.

#### Parallel Processor Operation

5       The four processors 100-103 shown in FIGURES 1 and 2 (abbreviated PP herein) perform most of the system's operations. The PP's each have a high degree of parallelism enabling them to perform the equivalent of many reduced instruction set computer (RISC)-like operations per cycle.  
10 Together they provide a formidable data processing capability, particularly for image and graphics processing.

Each PP can perform three accesses per cycle, through the crossbar switch to the memory, one for instructions and two for data. A multiply and an ALU operation can also be  
15 performed by each PP every cycle, as well as generating addresses for the next two data transfers. Efficient loop logic allows a zero cycle overhead for three nested loops. Special logic is included for handling logical ones, and the ALU is splittable for operating on packed pixels.

20       As discussed previously, to allow flexibility of use, the PPs can be configured to execute from the same instruction stream (Single Instruction Multiple Data (SIMD) mode) or from independent instruction streams (Multiple Instruction Multiple Data (MIMD) mode). MIMD mode provides  
25 the capability of running the PPs together in lock-step allotting for efficient synchronized data transfer between processors.

In order to relieve the programmer of the worries of accidental simultaneous access attempts of the same memory,  
30 contention prioritization logic is included in the crossbar, and retry logic is included in the PPs.

All the PPs 100-103 are logically identical in design, but there are two differences in their connections within the system. Firstly, each PP will be supplied with a  
35 unique hardwired two-bit identification number that allows

2025 RELEASE UNDER E.O. 14176

5

## Internal Interfaces

10

15

20

30

35

these sections of code by bounding them with LCK and ULCK instructions. Four signals 3020, one output by each PP, are routed between the PPs indicating when each is in this section of code. By testing these signals the PPs can  
5 execute code synchronously.

As mentioned above, global ports 3005 and local ports 3006 have signals 3210 and 3211 (FIGURE 32) to know when contention has occurred and a retry is required. When configured in SIMD mode, it is essential that all PPs pause  
10 instruction execution until all contentions have been resolved. There is thus a signal 3007 running between all PPs which is activated when any PP detects contention. The next instruction is only loaded by the PPs when this signal becomes inactive. This signal is also activated when the  
15 "master" SIMD PP 100 detects a cache-miss. In MIMD configuration signal 3007 is ignored.

In SIMD configuration stack coherency between the PPs must be maintained. When performing conditional calls, a signal 3008 is required therefore from the "master" SIMD PP  
20 100 to the "slave" SIMD PPs 101-103 that indicates that the condition was true and that the return address should be pushed by the "slave" PPs 101-103.

Another time when SIMD stack coherency must be maintained is when interrupts occur. In order to achieve  
25 this there is a signal 3009 which is activated by the "master" SIMD PP 100 which is observed by the "slave" PPs 101-103. All PPs 100-103 will execute the interrupt pseudo-instruction sequence when this signal is active.

Another SIMD interrupt-related signal 3010 indicates  
30 to the "master" PP 100 that a "slave" PP 101-103 has an enabled interrupt pending. This allows "slave" PPs 101-103 to indicate that something has gone wrong with a SIMD task, since "slave" PPs 101-103 shouldn't normally expect to be interrupted.

A number of interrupt signals 3011 are supplied to each PP. These allow a PP to be interrupted by any other PP for message-passing. Master processor 12 can similarly interrupt a PP for message-passing. The master processor  
 5 can also interrupt each PP in order to issue them with new tasks. In SIMD the interrupt logic in the "slave" PPs 101-103 must remain active for stack consistency and interrupts are handled slightly differently. This is discussed later.

10 The PP indicates with a signal 3012 to the transfer processor when a packet request is required. The transfer processor indicates when a packet request has been serviced with another signal 3013. In SIMD configuration only the "master" PP 100 will output packet requests to the Transfer  
 15 Processor.

#### Internal Structure

The bus structure of a PP is shown in FIGURE 30. There are three main units within the PP. These are the  
 20 program flow control unit 3002, the address unit 3001 and the data unit 3000. Each of these will now be discussed.

Program flow control (PFC) unit 3002 shown in FIGURE 31 contains the logic associated with the program counter 3100, i.e., the instruction cache control 3101, the loop  
 25 control 3102, the branch/call logic 3103 and the PP synchronization logic 3104. This logic controls the fetching of opcodes from the PP's instruction RAM 10-1, 10-5, 10-9 or 10-14. When a cache-miss occurs, it also communicates the segment address and the sub-segment number  
 30 to the transfer processor so that the code can be fetched.

Instruction pipeline 3105 is in the PFC Unit 3002. The PFC unit 3002 will therefore generate the signals 3112 necessary to control the address 3001 and data units 3000. The immediate data specified by certain opcodes are also

Interrupt enable 3107, interrupt flags 3106 and interrupt vector address generation logic is also in the PFC unit 3002. This prioritizes the active interrupts and injects a sequence of pseudo instructions into the pipeline 3105 to read the vector, save the program-counter 3100 and the status register 3108, and branch to the interrupt routine.

10 Packet request handshake signals 3012 and 3013 will  
also connect to the PFC unit 3002.

The PFC unit is the part of the PP whose behavior differs between PPs when configured in SIMD mode. The "master" SIMD PP 100 will behave more-or-less normally, but  
15 the "slave" PPs 101-103 will disable their cache logic 3018  
and flush the present flags 3109. Their loop logic 3102,  
synchronization logic 3104 and packet request signals 3012  
and 3013 are also disabled. The interrupt logic behavior  
is modified so that all PPs can behave identically.

20       Address unit 3001 shown in FIGURE 32 contains two identical subunits 3200 and 3201 each capable of generating a 16-bit byte address of a data location in the crossbar'd RAM 10. Within each subunit are four address registers 3202, four index registers 3203, four qualifier registers 25 3204, a modulo register 3205 and an ALU 3206. When two parallel data accesses are specified in the opcode, subunit 3200 outputs the address through global port 3005 and the other subunit (3201) through the local port 3006. When only one access is specified, then this address can come 30 from either subunit 3200 or 3201, unless a single common SIMD read is specified, in which case it is required to come from the "local" subunit 3201.

Address unit 3001 also supports retries if contention is detected on either, or both, global 3005 and local buses 35 3006.

5

10

15

20

30

registers 3300, but, as with immediates, the parallel moves cannot be specified in this case.

### Bus Structure

5       As can be seen from FIGURE 30, there are four buses 3014-3017 which run the length of the PP data path. These are used for all the data movement, and are a compromise between the number of buses (and read and write ports of registers) and the allowed sources and destinations for  
10 data unit operations.

      The left-most bus 3014 carries the 16-bit immediates (after left/right justification and sign-extension) to data unit 3000. This is also used to load immediates by passing them through ALU 3302 then out onto the register write bus  
15 3016.

      The next bus from the left 3015 carries any address unit 3001 or PFC unit 3002 register source to the data unit 3000. It is also used to carry the source data of stores going to memory 10 on global port 3005. It also carries  
20 the source of a register-to-register move occurring in parallel with an ALU operation.

      The next bus 3016 is used by loads from memory 10 on global port 3005 to any register, and by the results of a data unit operation to be written to any register. This  
25 bus carries a latch 3018 which is used temporarily for holding load data when the pipeline pauses through contention, synchronization or cache-misses.

      The right-most bus 3017 is used entirely by the Local port 3006 for loads and stores of data unit registers 3300  
30 from/to memory 10. This bus cannot access any registers other than the data unit's registers 3300. This bus carries a latch 3019 which is used temporarily for holding load data when the pipeline pauses through contention, synchronization or cache-misses.

### Pipeline Overview

The PPs' pipelines have three stages called fetch, address and execute. The behavior of each pipeline stage is summarized below:

5        **FETCH:** The address contained in program counter 3100 is compared with the segment registers 3110 and present flags 3109 and the instruction fetched if present. PC 3100 is post-incremented or reloaded from the loop start address 3111. If MIMD synchronization is active, then this  
10 allows/inhibits the instruction fetch.

**ADDRESS:** If the instruction calls for one or two memory accesses, then the address unit 3001 will generate the required address(es) during this stage. The five most-significant bits of the address(es) are supplied to  
15 crossbar 20 for contention detection/prioritization.

**EXECUTING:** All register-to-register data unit 3000 operations and any other data movements occur during this stage. The remaining 11 bits of crossbar address(es) are output to the RAMs 10 and the data transfer(s) performed.  
20 If contention is detected, then this stage is repeated until it is resolved. If the PC 3109 is specified as a destination (i.e., a branch, call or return) then the PC 3100 is written to during this stage, thus creating a delay slot of two instructions.

25

### **MEMORY**

      Each PP accesses three separate memory spaces,  
      64M bytes of *off-chip* word-aligned code space. (From  
      on-chip cache).  
30        64K bytes of on-chip crossbar'd memory 10. This is  
      referred to as data space.  
      64K bytes of on-chip I/O space in which resides the  
      Parameter RAMs, the message registers and the  
      semaphore flags.



The I/O spaces for each PP 100-103 are isolated from each other so that code need not calculate addresses unique to each PP when accessing I/O space. Thus each PP sees its own parameter RAM at the same logical address. The same  
 5 applies for the message registers and semaphore flags. The master processor, however, can uniquely address each PP's I/O space.

The 64K bytes of memory is for one embodiment only and any expansion or modification can be made thereto.

10

#### Program Flow Control Unit

The logic within program flow control unit 3002, (FIGURE 31), works predominantly during the fetch pipeline stage, affecting the loading of the instruction pipeline.  
 15 However since the instruction pipeline is resident in the PFC unit 3002, it must also issue signals 3112 to the address 3001 and data units 3000 during the address and execute pipeline stages. It also receives signals from address unit 3001 that indicate when contention has  
 20 occurred, thus pausing the pipeline.

#### Cache Control

The 512-instruction cache has four segments, each with four sub-segments. Each sub-segment therefore contains 32  
 25 2instructions. There is one present flag 3109 for each sub-segment. Since program counter 3100 is 24 bits, the segment registers 3110 are each 17 bits. The instruction opcodes are 32-bits wide.

The 9-bit word address used to access the instruction  
 30 RAM is derived from the least-significant 7 bits of program counter 3100 and two bits from the segment address compare logic 3113. This compare logic must work quickly so as to avoid significantly delaying the RAM access.

If the most-significant 17 bits of program counter  
 35 3100 are not matched against one of the segment address

registers 3110, then a segment-miss has occurred. The least recently used segment is chosen to be trashed by logic 3114, and its sub-segment present flags 3109 are cleared. If, however, the most-significant 17 bits of the  
 5 Program Counter 3100 are matched against one of the segment address registers 3110 but the corresponding sub-segment flag 3109 is not set, then a sub-segment miss has occurred.

If either type of cache-miss occurs, the pipeline is paused, and a cache-miss signal 3115 sent to transfer  
 10 processor 11. When a cache-miss acknowledge signal is supplied by the TP 11, the most-significant 17 bits of the PC 3100, and the 4 bits representing the sub-segment to be filled are output onto the TP's bus. (This requires a crossbar connection 0-3, 0-8, 0-13 or 0-18 between the PP's  
 15 instruction bus, horizontal 7, and the TP's bus, horizontal 0). The TP's acknowledge signal 3115 is then deactivated. When the sub-segment has been filled by TP 11, a cache-filled signal 3115 is sent to the PP which causes the appropriate sub-segment present flag 3109 to be set,  
 20 deactivates the PP's cache-miss signal 3115, and instruction execution recommences.

If the PP is interrupted at any time while waiting for a cache-miss request to be serviced, the cache miss service is aborted. This prevents needless fetches of unwanted  
 25 code.

In SIMD configuration the present flags 3109 of the "slave" PP's 101-103 will be held cleared and the cache logic 3101 ignored. The "slave" PP's 101-103 will load instructions (supplied by the "master" PP 100) into their  
 30 pipeline whenever the SIMD pause signal 3007 is inactive. The "master" PP's cache 3101 behaves normally, but it too will pause its pipeline whenever the SIMD pause signal 3007 is active. (Such a condition will occur if one of the "slave" PPs 101-103 detects contention). In MIMD



loop end address register 3116-3118 should be compared with the PC 3100. These CLD bits will be cleared to zero (no loops active) by reset, and by interrupts once the SR 3108 has been pushed. Loops are numbered 1 to 3 with 1 being the outer-most loop. The user must set the MLD and CLD bits to the desired values in order to activate the loop logic. When all loops have been completed the CLD bits will be zero.

Since the CLD bits are automatically decremented by the loop logic during the fetch pipeline stage, the status register 3108 should not be written to during the last two instructions within a loop.

Once the loop logic 3102 has been activated (by a non-zero value in the CLD bits) the 16-bit loop end address register (one of 3116-3118) indicated by the CLD bits is compared during each instruction fetch with the 16 least-significant bits of the unincremented PC 3100. If they are equal and the associated loop counter (one of 3119-3121) is not 1, then the loop start address register 3111 contents are loaded in the PC 3100, the loop counter (one of 3119-3121) is decremented and the MLD bits are copied into the CLD bits.

If, however, the unincremented PC 3100 and loop end address register (one of 3116-3118) are equal and the relevant loop counter (one of 3119-3121) is 1, then the CLD bits are decremented by 1, the relevant loop counter (one of 3119-3121) is reloaded from its associated loop reload register (one of 3122-3124), and the PC 3100 increments to the next instruction.

Since the loop end address registers 3116-3118 are only 16-bits, this means that loops cannot be more than 64K instructions long. Care should also be taken if branching or calling out of loops as the 16-bit value of the currently in-use loop end address register (one of 3116-3118) may be encountered accidentally. Users should set

the CLD bits to zero before attempting this to be certain of not having a problem. Loop end address compare is disabled during the two delay slot instructions of a branch or call in order to help returns from interrupts.

5        Since the loop logic operates during the fetch pipeline stage it is possible to combine looping with MIMD synchronization if desired. Interrupting loops is similarly not a problem. Looping in SIMD is controlled by the "master" 100 SIMD PP's loop logic. The "slave" PPs'  
10 101-103 loop logic can still operate since their program counters 3100 are ignored.

There are various permutations on the above arrangement which can be used. A slightly more user friendly method is to have three 24-bit loop end registers  
15 with comparators, and three 24-bit loop start address registers. Each loop would be enabled by a single bit in the status register.

When executing MIMD programs that are working on a common task, there is usually the need to communicate  
20 between processors. The system supports both message-passing and semaphores for "loose" communication, but when executing tightly-coupled programs, the need to exchange information on a cycle-by-cycle basis is required. This is where synchronized execution is of benefit.

25        Within each PP's SYNC/PP# 3104, register there are four bits one relating to each PP. The other PP(s) to which a particular PP will synchronize is indicated by writing a one to the bits corresponding to those PP(s). The other PP(s) which are expecting to be synchronized will  
30 similarly have set the appropriate bits in their SYNC/PP# 3104 register(s).

Code that is desired to be executed in synchronization is indicated by bounding it with LCK (Lock) and ULCK (Unlock) instructions. The instructions following the LCK,  
35 and those up to and including the ULCK, will be executed in

TOP SECRET

lock-step with the other PP(s). There must therefore be the same number of instructions between the LCK and ULCK instructions in each synchronized PP.

The knowledge that synchronized code is being executed  
 5 is recorded by the "S" (synchronized) bit (26) in status register 3108. This bit is not set or reset until the master phase of the address pipeline stage of the LCK or ULCK instructions respectively, but the effect of the LCK or ULCK instruction affects the fetch of the next  
 10 instruction during the slave phase. This bit (26) is cleared by reset and by interrupts, once the status register 3108 has been pushed.

When a PP encounters a LCK instruction (decoded during the slave phase of the address pipeline stage) it will  
 15 output a signal 40 to the other PPs 100-103 saying that it is executing a piece of synchronized code. It will then AND the incoming sync signals from the other PPs with which it is desiring to be synchronized, and only when all those processors are outputting sync signals 40 will the next  
 20 instruction be fetched into the pipeline. This will occur coincidentally in all the synchronized PPs because they too will not proceed until the same set of matching sync signals are active. It is therefore possible to have two different synchronized MIMD tasks running concurrently,  
 25 because each will ignore the sync signals of the other.

Since it is the instruction fetch that is synchronized, it is possible to interrupt synchronized code. This will immediately cause the PP's sync signals 40 to become inactive. Cache-misses and contention will have  
 30 a similar effect, keeping the machines in-step. In the case of contention, however, the instruction following the one experiencing contention will have already been fetched into the pipeline before the pipeline pauses.

It is possible to put IDLE instructions into  
 35 synchronized code, thus holding the operation of all the



experiences a cache-miss when fetching instruction 4, but since the PPs are now unlocked PP2 carried an unimpeded.

It should be noted that this simple example can be further simplified by combining instructions 0 with 1, and  
 5 2 with 3. (i.e., LCK11 ST followed by ULCK11 LD). This way just the loads are synchronized, but that is all that is required in this case.

Synchronization in SIMD is implicit, so the LCK and ULCK instructions have no purpose and so have no effect if  
 10 coded. The S bit (26) in the Status Register 3108 will have no effect if a program should set it to one.

### Interrupts and Returns

Interrupts must be locked-out during the two delay  
 15 slots after the PC 3100 has been loaded. This prevents having to save both the current PC 3100 value, and the branch address, and restore them on the return. Loads of the PC 3100 are forbidden during delay slot instructions, but if a user somehow does this, then the lock-out period  
 20 isn't extended; otherwise, it would be possible to lock-out interrupts indefinitely.

Like many processors, there is a global interrupt enable bit (27) (I) in status register 3108. This can be set/reset by the user to enable/disable all interrupts,  
 25 except the master task interrupt, and the illop interrupt. Bit (27) is cleared by reset and by the interrupt pseudo-instructions after status register 3108 has been pushed.

Returns from interrupts are executed by the sequence POP SR, POP PC, DELAY1, DELAY2. The I (27), S (26) and CLD  
 30 (23) and (22) bits of status register 3108 are loaded by the POP SR before the DELAY2 instruction, but their effects are inhibited until the branch (POP PC) instruction has completed. This prevents them becoming effective before the return has completed.



There is provision for up to 16 interrupt sources on each PP 100-103. Of these, eleven are assigned, the others are left for future expansion. Those assigned are:

5	Master Task	The master processor wishes the PP(s) 100-103 to run a new task. (Always enabled)
	Illop	An illegal opcode was detected. (Always enabled)
10	SIMD error	Applicable only to the "master" SIMD PP 100. It is an OR of all enabled interrupts of the three "slave" PPs 101-103.
	Illadd	A non-existent on-chip address was accessed.
	Contention	Contention was detected. Interrupt is taken after contention is resolved in the normal manner.
15	Packet Request	The transfer processor has exhausted the PP's packet request linked-list.
	Master Message	Occurs when the master processor 12 writes to the PP's message register.
20	PP0 Message	Occurs when PP0 writes to the PP's message register.
	PP1 Message	Occurs when PP1 writes to the PP's message register.
	PP2 Message	Occurs when PP2 writes to the PP's message register.
25	PP3 Message	Occurs when PP3 writes to the PP's message register.

#### Interrupt Registers

There are two registers that control interrupts; the interrupt flag register 3106 (INTFLG), and the interrupt enable register 3107 (INTEN).

Interrupt enable register 3107 has individual enable bits for each interrupt, except for the master task and illop interrupts which have their associated enable bits hard-wired to one. This register is cleared to all zeros

(except the two wired to one) by reset. Bits 15 to 0 are unimplemented.

Interrupt flag register 3106 has an individual flag for each interrupt source. This flag is latched by the source signals which are each active for a single cycle. This register is cleared to all zeros by reset. Bits 15 to 0 are unimplemented. Those marked as reserved will also be hardwired to zero. Any flag can be cleared by writing a 1 to it. Writing a zero has no effect. This allows the flags to be polled and cleared by software if desired instead of generating interrupts. When an interrupt is taken, the associated flag will be cleared automatically by the hardware. If a flag is being set by a source at the same time as it is being cleared, then the set will dominate.

Interrupt flag register 3106 can be written with ones and zeros like a normal data register once the R (restore registers) bit (19) of status register 3108 is set. This allows task state restoring routines to restore the interrupt state.

When interrupts are enabled, by setting the I bit (27) in status register 3108, the interrupts are prioritized. Any enabled interrupt whose flag becomes set will be prioritized, and an interrupt generated at the next possible opportunity. A sequence of three pseudo-instructions is generated which

1. generates the address of interrupt vector and fetches it into the PC 3100, having first copied the PC into RET 3103, and clears the interrupt flag in 3106 unless it is being simultaneously set again;
2. Pushes RET 3103; and
3. Pushes SR 3108 and clears the S (26), I (27) and CLD (22) and (23) bits in SR 3108. It also

disables the functions associated with these bits until the execute stage has completed.

Contention resolution must be supported by the above sequence, so it may take more than three cycles to execute.

- 5 Similarly a cache-miss on either of the first two instructions of the interrupt routine will cause the pipeline to pause.

- The interrupt vectors are fetched from the PPs' own Parameter RAM 10. Since these exist at the same logical address for each PP 100-103, the interrupt logic in each PP 100-103 generates the same vector addresses.

- It is a consequence of the pipelining that neither of the first two instructions of an interrupt routine can be a LCK instruction. For similar reasons the interrupt logic must disable interrupts 3106, sync logic 3104 and loop logic 3102 until the execute stage of the third pseudo-instruction has completed. This prevents these functions from being active during the fetching of the first two instructions of the interrupt routine.

- 20 Interrupts are handled slightly differently in SIMD from MIMD. In order to maintain stack coherency there is a signal from the "master" PP 100 to the "slave" PPs 101-103 that indicates that it is taking an interrupt. This causes the "slave" PPs 101-103 to execute their sequence of interrupt pseudo-instructions. It really doesn't matter which interrupt vector they fetch since their PCs 3100 are ignored anyway.

- In SIMD configuration there is also the need to pass back to the "master" PP 100 the fact that a "slave" PP 101-103 has detected an enabled interrupted event. This could be contention, or an illegal address access or a message interrupt. Since any one of these is almost certainly an error they are handled by only one interrupt level on the "master" PP 100. There is one signal 3010 running from the "slave" PPs 101-103 to the "master" PP 100 which is the

5        If an interrupt occurs (from the "master" PP 100)  
while the SIMD pause signal 3007 is active, the issuing of  
the "master" to "slaves" interrupt signal 3009 will be  
delayed until the cause of the pause has been removed. If  
the cause of the pause is a cache-miss, the cache-miss will  
10 be aborted and the interrupt can be taken immediately.

Branches and calls are achieved by writing into the PC 3100, which is an addressable register like any other PP register at the same time that the branch address is written into the PC 3100 the value of PC+1 is copied into the return address register, RET 3103. This is the value required for a return if the branch is really a call. This RET register 3103 is then programmed to be pushed onto the stack by either of the delay slot instructions in order to make it into a call. To allow conditional calls there is an instruction for conditionally pushing the return address. This only occurs if the branch is taken.

As described earlier, instructions specifying the PC 3100 as the destination will lock-out interrupts until after the second delay instruction has been fetched. This prevents problems with the branch address and/or return address getting lost. During this period synchronization is also disabled as described earlier. In order to prevent 30 problems on returns from interrupts with loop logic 3102 becoming activated too early, loop end address compare is also disabled during the two delay slot instructions.

### Status Register

Status Register 3108 is resident in the PFC unit 3002 and shown in FIGURE 34. Each bit's function is described in the following sections.

5        The N - Negative bit (31) is set by certain instructions when the result was negative. Writing to this bit in software will override the normal negative result setting mechanism.

The C - Carry bit (30) is set by certain instructions  
10 when a carry has occurred. Writing to this bit in software  
will override the normal result carry setting mechanism.

The V - Overflow bit (29) is set by certain instructions when an overflow has occurred. It is not a permanently latched overflow. Its value will only be maintained until the next instruction that will set/reset it is executed. Writing to this bit in software will override the normal result overflow setting mechanism.

The Z - Zero bit (28) is set by certain instructions when the result was zero. Writing to this bit in software will override the normal zero result setting mechanism.

The I - Interrupt Enable bit (27), which is set to zero by reset and interrupts, is a global interrupt enable. It enables all the interrupts whose interrupt enable bits are set. Due to normal pipeline delays changing the value of this bit will have no effect until after the execute stage has completed.

The S - Synchronized code execution bit (26), which is set to zero by reset and interrupts, indicates that synchronous MIMD code execution is operating. Instructions will only be fetched when all the PPs indicated by the SYNC bits in the SYNC/PP# 3104 register are outputting active sync signals 40. This bit's value is ignored in SIMD configuration.

The MLD - Maximum looping depth bits (24) and (25),  
35 which are set to zero by reset, indicate how many levels of

loop logic are operating. 00 - indicates no looping, 01 - just loop 1, 10 - loops 1 and 2, 11 - all three loops active.

The CLD - Current looping depth bits (22) and (23),  
 5 which are set to zero by reset, indicate which of the Loop End registers is currently being compared against the PC. 00 - indicates no looping, 01 - Loop End 1, 10 - Loop End 2, 11 - Loop end 3. These bits are set to zero by reset and by interrupts once status register 3108 has been  
 10 pushed.

The R - Restoring registers bit (19), which is set to zero by reset, is used when restoring the state of the machine after a task switch. When set to a one, it allows interrupt flag register 3106 to be written with ones and  
 15 zeros like a normal register, and also the message registers to be restored without causing new message interrupts. It also enables the Q bit (17) of status register 3108 to be written to for similar reasons. The R bit (19) will therefore only be used by task restoring  
 20 routines.

The U - Upgrade packet request priority bit (18), which is set to zero by reset, is used to raise normal background priority packet requests to foreground. Its value is transmitted to transfer processor 11 and is used  
 25 in conjunction with the Q bit's value to determine the priority of transfer requests. This bit remains set until reset by software.

The Q - Queued packet request bit (17), which is set to zero by reset, indicates that the PP has a packet  
 30 request queued. It becomes set one cycle after the P bit (16) of the status register 3108 is written with a one. This bit's value is transmitted to transfer processor 11 and used in conjunction with the U bit's (18) value to determine the priority of transfer requests. This bit is  
 35 cleared by transfer processor 11 once the PP's linked-list

of packet requests has been exhausted. If this bit is being set (via the P bit (16)) by software at the same time as transfer processor 11 is trying to clear it, then the set will dominate. Writing to this bit directly has no effect, unless the R bit (19) in status register 3108 is set, when this bit can be written with a one or zero. This can be used to de-queue unwanted packet requests, but is more normally needed for restoring interrupted tasks.

The P - Packet Request bit (16), which is set to zero by reset, is a one-shot single-cycle bit, used to set the Q bit (17) in status register 3108. This initiates a packet request to transfer processor 11. The P/Q bit mechanism is to allow read-modify-write operations on status register 3108 without accidentally initiating packet requests if the packet request bit was cleared by the TP 11 between the read and write.

All unimplemented status register bits 3108 will read as zero. Writing to them has no effect. They should only be written with zeros to maintain future device compatibility.

#### Synchronization Indicators

The four SYNC bits, which are set to zero by reset, are used to indicate to which PP a MIMD PP wishes to synchronize. When executing code bounded by LCK and ULCK instructions, instruction fetches will not proceed unless all those processors indicated by one in the corresponding SYNC bits are outputting sync signals 40. These bit values are ignored in SIMD configuration.

The two PP# bits are unique to each PP 100-103. They are hardwired to allow software to determine which PP it is running on, and thus calculate correct unique addresses. Writing to these bits has no effect.

The coding of these bits is; 00 - PP0 100, 01 - PP1 101, 10 - PP2 102 and 11 - PP3 103. PP0 100 is the

"master" SIMD PP. The associated start addresses of the PPs' Local crossbar RAMs are; 0000h - PP0 100, 2000h - PP1 101, 4000h - P2 102 and 6000h - PP3 103.

Pipeline control can be difficult. The reason for  
5 this is the number of concurrent operations that inter-  
relate as demonstrated below:

Instruction fetch with associated cache management.  
Address generations with various addressing modes.  
Crossbar access requests with independent contention  
10 resolution.  
Memory transfers.  
Loop address compare, with PC load/increment.  
Loop count decrement/reload.  
Looping depth count decrement/reload.  
15 Multiply.  
Shift.  
Add/subtract.  
Synchronization with other PPs.  
Interrupt detection/prioritization.

The pipeline "events" that cause an "abnormality" in the straightforward execution of linear code are:

```

25      Instruction cache-miss
      Contention on the Global and/or Local buses
      Loops
      Branches and calls
      Interrupts
      Idling
      Synchronization

```

30 In the following sections the events are shown  
diagrammatically. The abbreviations "pc+1" and "pc"  
indicate whether the Program Counter 3100 is incremented  
normally, or not, respectively. The pipeline boundaries  
marked are the Stages, which consist of the slave clock  
35 phase followed by the master clock phase, ie. | s:m |.  
Where cycles may be repeated an indefinite number of times  
this is shown by "| ... |".

### Cache-miss Pipeline Sequence

40        The pipeline sequence for a cache-miss is shown in  
FIGURE 35. The cache-miss is detected during the slave



5  
10

15

20

25

30

### Contention Resolution Pipeline Sequence

35

5

10

20

25

30

### Loop Control Pipeline Sequence

35

Loop End 1 3116, Loop Count 1 3119 and Loop Reload 1 3122 registers). It contains 2 instructions, and the loop counter value before starting the loop is 2. The principles can be extended to all three loops.

5        In this example, when PC 3100 is found (during the slave phase) to be equal to loop end register 3116, loop counter 3119 is compared to 1. As it is not equal, the PC 3100 is reloaded from start address register 3111, loop counter 3119 is decremented by 1 and the current looping  
10    depth bits 3108 (bits (22) and (23)) are reloaded from the maximum looping depth bits 3108 (bits (24) and (25)) (in this example the CLD bits values don't change).

      The loop is repeated again, but this time when the end of loop is detected, loop counter 3119 is 1, so PC 3100 is  
15    incremented to the next instruction instead of being loaded from start address register 3111. Loop counter 3119 is reloaded from loop reload register 3122 and current looping depth bits 3108 (bits (22) and (23)) are decremented by 1.

      The pipeline sequence for a branch or call is shown in  
20    FIGURE 38. When the branch address is written into the PC 3100 the value of PC+1 (calculated during the slave phase) is loaded into RET 3103. This is the address of the instruction after the second delay instruction, and is the return address for a call.

25        The branch address can come from memory, a register, an immediate 24-bit value or by adding a 24-bit index to the current PC value in 3100.

      Difficulties with saving the branch address and the return address would occur if interrupts were allowed  
30    during the delay slot instructions. In order to prevent this interrupts are locked out during the fetch pipeline stage of the two delay slot instructions. This requires decoding a PC 3100 destination during the slave phase of the address pipeline stage. Lockout of interrupts will  
35    occur with conditional branches, as the condition isn't

testable until after the two delay slot instructions have been fetched.

As described in the synchronization section, branches and calls are treated as one instruction as far as synchronization is concerned. Thus the PP's sync signal 40 goes inactive during the two delay slot instructions, with the timing shown. This is also true for conditional branches and calls regardless of the condition.

Also, since conditional calling is done by pushing RET 3103 (return address) only if the conditional branch is taken, then there is a potential problem with conditional calls in SIMD, since the "slave" PPs 101-103 don't know if the branch was taken. They therefore wouldn't know if they should push RET 3103, and thus could lead to stack inconsistency. In order to fix this the signal "SIMD branch-taken" 3008 is output from the "master" SIMD PP 100 to the "slave" PPs 101-103 which they use to determine if their PRET instructions should push RET 3103. This is taken active (or left inactive) with the timing shown.

### Interrupts

The pipeline sequence for an interrupt is shown in FIGURE 39. The sequence is that for any machine in MIMD or SIMD, but if the interrupt source is a "slave" SIMD PP 101-103 then the sequence is kicked off by the "slave" PP to "master" PP interrupt signal 3010 as shown. The "slave" PP 101-103 will wait for the "master" PP 100 to output the "master" PP to "slave" PPs interrupt signal 3009 as shown.

Once an enabled interrupt is detected, the sequence of pseudo-instructions is commenced. The first instruction calculates the interrupt vector address and fetches the vector into the PC 3100 and copies the old PC value (return address) into RET 3103. The second instruction pushes RET 3103. The third instruction pushes the SR 3108 and clears its S, I and CLD bits.

Note that the first two instructions of the interrupt routine are being fetched before the SR 3108 has been pushed and its S, I and CLD bits cleared. The functions of the S, I and CLD bits are thus disabled by the interrupt logic until SR 3108 has been pushed, and the S, I and CLD bits cleared.

#### IDLE Pipeline Sequence

The pipeline sequence for an IDLE instruction is shown in FIGURE 40. The IDLE instruction is decoded before the end of the slave phase of its address pipeline stage, allowing it to stop the PC 3100 from being incremented and the pipeline from being loaded with the next instruction. The MIMD pause is taken inactive and the SIMD pause signal is activated. Instruction fetching halts until the interrupt logic detects an enabled interrupt. This will kick off the sequence of pseudo-instructions once an enabled interrupt is detected. If the interrupt source comes from a "slave" SIMD PP 101-103, then the interrupt sequence isn't kicked off until the "master" PP to "slave" PPs interrupt signal 3009 is activated.

If parallel transfers are coded with an IDLE instruction they will occur when the interrupt occurs, before the interrupt routine is executed.

#### Synchronization

The pipeline sequence for a synchronized MIMD or SIMD PP waiting for an incoming sync signal to become valid is shown in FIGURE 41. The next instruction is not fetched into the instruction pipe until all the desired PPs are outputting active sync signals.

#### Address Unit

The logic within address unit 3001 works predominantly during the address pipeline stage, calculating the

10

15

20

30

35

cannot be referenced by the following instruction. They may be referenced by the next-but-one instruction. This allows interrupts to occur without undesired consequences.

The global and local subunits 3200 and 3201 are identical apart from the register numbers, so one description will serve for both. There are however slight differences in how the two units are connected and used which will be highlighted, but the internal content of the subunits is the same.

10        Within each subunit are four 16-bit address registers 3202 (A0-A3) or (A4-A7). These contain indirect addresses which are either used unchanged or to which indices are added. If an index is added, then there is the option of replacing the previous address register value in 3202 by  
15        the value created by indexing.

The values within the address registers 3202 are always interpreted as byte addresses, regardless of the data size being transferred. Non-aligned word or half-word transfers can be specifically coded but this requires two  
20        instructions. This is discussed later.

All address accesses of the PPs 100-103 must be sourced from an address register 3202 or 3222. The capability of coding an immediate address within the opcode is not provided. This is considered to be of low  
25        significance since SIMD tasks would not normally wish to specify the same address for each PP. It is also thought to be of low importance for MIMD since MIMD algorithms should be written in such a manner that they can be run on any PP.

30        Address register A7 3227 is reserved as the stack pointer. It can be referenced like any other address register 3202 or 3222, but obviously care must be taken if adjusting A7's value, as interrupts can occur at any time. PUSH, POP and interrupts treat pushes as pre-decrement, and  
35        pops as post-increment.

Within each subunit 3200 or 3201 are four 16-bit index registers (X0-X3) 3203 and (X4-X7) 3223. The contents of these can be requested by the opcode to be added to, or subtracted from, the contents of the specified address register 3202 or 3222, in order to perform indexed addressing. This addition/subtraction can be performed either before or after the address is put out onto crossbar 20, thus allowing pre- or post-indexing respectively. The address created by pre-indexing can optionally be stored back into address register 3202 or 3222. This is compulsory for post-indexing.

If only one access is specified by the opcode, then any one of the four index 3203 or 3223 registers within the same subunit as the address register 3202 or 3222 can be specified as the index source, (eg. A0 and X2, A6 and X4, ...). The indexing modes that can be specified are pre- or post-, addition or subtract, with or without address register 3202 or 3222 modify.

If two parallel accesses are specified, then the index register 3203 or 3223 with the same suffix as the address register 3202 or 3222 is used (eg. A2 and X2, A5 and X5), and only post-addition-indexing is available.

The values contained within the index registers 3203 and 3223 are always interpreted as byte addresses, regardless of the data size being transferred.

An alternative indexing method to index register indexing is short-immediate or implied immediate indexing. Short-immediate indexing, which is available when only one access is specified, allows a 3-bit short immediate value to be used as the index. As with index register indexing this can be either pre- or post-, addition or subtraction, with or without address register 3202 or 3222 modify.

If two parallel access are coded then only an implied immediate of +1 with post-indexing, and -1 with pre-indexing, can be specified. These allow stacks of 8, 16 or



32 bits to be accessed even when two parallel transfers are coded.

When specifying short-immediate or implied immediate, the immediate value is shifted 0, 1 or 2 bits left by shifter 3208 or 3228 if the specified word size is 8, 16 or 32 bits, respectively, before being added to the value from address register 3202 or 3222. The short-immediate index is thus 0-7 "units", and the implied immediate is +/-1 "unit", where a "unit" is the data size. The address register is not shifted as it always contains a byte address.

Associated with each address register (A0-A3) 3202 or (A4-A7) 3222 is an 8-bit address qualifier register (Q0-Q3) 3206 or (Q4-Q7) 3224. These qualifier registers contain extra information required for the access which cannot be fitted into the opcode. This information typically isn't required to be modified on a cycle-by-cycle basis.

Since A7 3227 is assigned to be the stack pointer, bits 6-0 of Q7 3229 are hardwired to 0000010 respectively. The individual bit functions of the Q registers 3204 and 3224 are described below:

A PP's address space is divided into two halves; data space (the crossbar'd memory 10) and I/O space (the parameter RAMs, message registers and semaphore flags). If this bit is a 1, then the access is performed to the I/O space. 0 directs the access to the crossbar'd RAM 10.

If the power-of-2 modulo bit has value 1, then it indicates the desire to break the carry path on the address adder 3206 or 3226 at the position indicated by a 1 (or perhaps several 1s) in the modulo register, M0 3205 or M4 3225, associated with the subunit 3200 or 3201. This allows power-of-2 dimension matrix addressing to be performed. If this bit is 0 then the address adder 3206 or 3226 behaves as a normal 16-bit adder/subtractor.

If the reverse-carry addressing bit is set to a 1, then reverse-carry addressing is enabled. This causes the carry path of the address adder/subtracter 3206 or 3226 to reverse its direction. When specifying indexed addressing  
 5 with a power-of-2 index (eg. 8, 16, 32 etc.) this has the effect of counting in a manner required by FFTs and DCTs. If this bit is 0 then the address adder 3206 or 3226 behaves as a normal 16-bit adder/subtracter.

The common SIMD load bit when set to 1 specifies that  
 10 if a load is specified, then it should be a common SIMD load. This bit, due to the nature of the common SIMD load, is only relevant to Q4-Q6 3224 of the "master" SIMD PP 100 when specifying a load. This will cause local buses 3006 of the PPs to be series connected for the duration of the  
 15 load. If this bit is zero, then the common SIMD load function will be disabled. Setting this bit in "slave" PPs 100-103, or other than Q4-Q6 of the "master" SIMD PP, will have no effect. Stores are unaffected by this bit value.

When the sign extend bit is set to a 1, loads of half-  
 20 words or bytes will have bit 15 or bit 7, respectively, copied to all the most-significant bits when loaded into the PP register. This is a function of the aligner/extractor. If this bit is a 0, then all the most-significant bits will be zero-filled.

25 The two size bits specify the size of the data to be transferred. The codings are 00 - 8 bits, 01 - 16 bits, 10 - 32 bits, 11 - reserved. These bits control the function of the aligner/extractor 3003, the byte strobes on stores, and the sign extend function.

30 Address ALUs 3206 and 3226 are normal 16-bit adder/subtracters except they can have the direction of their carry paths reversed or broken.

When performing in-place FFTs the addresses of either the source data or the results are scrambled in a way that  
 35 make them difficult to access. The scrambling however has

an order to it that allows fairly easy unscrambling if the direction of the carry path of address adder 3206 or 3226 is reversed. This feature which is common on DSPs is usually referred to as reverse-carry addressing, or bit-reversed addressing.

A power-of-2 index (eg. 8, 16, 32 ...) equal to the power-of-2 number of points in the FFT divided by 2, is added onto the address from the address register 3202 or 3222 using a reversed carry ripple path. The resulting value is used as the address and stored in the address register 3202 or 3222. This produces the sequence of addresses required to unscramble the data, e.g., if the index is 8 and the initial address register value was 0, then the sequence 0, 8, 4 C, 2, A, 6, E, 1, 9, 5, D, 3, B, 7, F is produced.

The reverse-carry feature will operate with any indices other than power-of-2 numbers, but may not yield any useful results. This feature is only operative when the reverse-carry bit in Q register 3204 or 3224 associated with the specified A register is set to 1.

When distributing data around crossbar memories 10, there may well be situations where a "wrap-around" is required in a particular dimension, in order to access consecutive data, handle boundary conditions or address arrayed data. In order to easily support this, the ability to break the carry path of address adder 3206 or 3226 at one or more chosen places is provided.

The location of the break(s) is determined by modulo register M0 3205 or M4 3225. A 1 located in bit n of a Modulo register will break the carry path between bits n-1 and n of the address adder. This allows a  $2^n$  modulo buffer to be implemented. Any number of 1s can be programmed into the Modulo registers 3205 or 3225 as desired. This allows multi-dimensional arrays to be implemented, with each dimension being a power-of-2 modulo amount.

This feature is only active when the power-of-2 modulo bit in qualifier register 3204 or 3224 associated with specified address register 3202 or 3222 is set to 1. Otherwise normal linear addressing applies.

5

#### Local and Global Ports

The main feature of global ports 3005 and local ports 3006 are the aligner/extractors 3003. They handle the movement of 8, 16 and 32-bit data, sign-extension, non-aligned access and common SIMD loads. To achieve these functions the aligner/extractors 3003 are basically a collection of byte multiplexers, wired to give the required operations. Each global port 3005 or local port 3006 operates independently, so a description of one applies for the other. Common SIMD load is the exception to this statement and is discussed with the other functions below:

The data size of a load or store is defined within qualifier register 3204 or 3224 associated with the specified address register 3202 or 3222. Valid options are 8, 16 or 32 bits. The data size can thus vary on a cycle-by-cycle basis, dependant upon which address register 3202 or 3222 is accessed and the values within its qualifier register 3204 or 3224.

A full 32-bit word of data is always transferred across the crossbar between memory 10 and the PP 100-103, or vice-versa, even when the specified word size is 8 or 16 bits. When performing loads of 8 or 16-bit quantities, the appropriate byte(s) are extracted from the 32-bit word according to the LS bits of the address and the word size. This is right-shifted if required, to right-justify the data into the PP register destination. The upper bytes are filled either with all zeros, or if sign extension was specified in the qualifier register 3204 or 3222, the MS bit (either 15 or 7) is copied into the most significant bytes.

When storing 8 or 16-bit quantities to the crossbar'd memory 10 the (right-justified) data is repeated 4 or 2 times, respectively, by aligner/extractor 3003, to create a 32-bit word. This is then written across crossbar 20 accompanied by four byte strobes which are set according to the LS bits of the address and the data size. The appropriate byte(s) are then written into the memory.

The above description of data loads and stores assumes that the addresses are aligned. That is, 16-bit accesses 10 are performed to/from addresses with the LS bit = 0, and 32-bit accesses are performed to/from addresses with the two LS bits - 00. (8-bit quantities are always aligned).

Provision however is made to allow accesses of non-aligned 16 or 32-bit data. This is not automatic, but 15 requires the user to specifically encode loads or stores of the upper and lower parts of the data separately. There are thus four instructions available that "load upper", "load lower", "store upper" and "store lower" parts of the data. These instructions use the byte address and data 20 size to control aligner/extractor 3003 and, in the case of loads, only load the appropriate part of the destination register. This requires the registers to have individual byte write signals. For this reason non-aligned loads will be restricted to data registers 3200 only.

25 In practice the "load lower" and "store lower" instructions are the normal load and store instructions. If the address is aligned then the transfer is completed by the one instruction. If followed (or preceded) with the "upper" equivalent operation, then nothing will be 30 transferred. If the address is not aligned, then only the appropriate byte(s) will be stored to memory or loaded into a register.

Some examples of non-aligned operation may help the explanation here and are shown in FIGURES 42 and 43. These 35 are all little-endian examples which are self explanatory.

Common SIMD Load

There is sometimes the need, such as in convolution, to perform two accesses in parallel in each machine each cycle. One of these is data coming from anywhere in the crossbar'd memory 10 via global ports 3005, and the other is information "common" to each PP 100-103, such as a kernel value. This would therefore be entering via local port 3006. In order to pass this information to all local ports 3006 simultaneously from one source of data, there are unidirectional buffers that series-connect local crossbar data buses 6.

These series connections are only made in SIMD, when an address register 3222 in the local address subunit 3201 is accessed with the common SIMD load bit set in its associated qualifier register 3224, and a load is specified. Under all other conditions local data buses 6 are disconnected from each other. When the series connections are made, the addresses output by PPs 1-3 101-103 (the "slave" SIMD PPs) are ignored by the crossbar 20.

Since the series connecting buffers are unidirectional, the common data can only be stored in the four crossbar RAMS 10.0, 10.2, 10.3 and 10.6 opposite the "master" SIMD PP, PP0 100. (ie. in the address range 0000h - 1FFFh).

Contention Resolution

The purpose of contention resolution is to allow the user to be freed from the worries of accidentally (or deliberately) coding two simultaneous accesses into the same RAM by any two devices in the system. There are seven buses connected to each crossbar RAM. It would therefore be a considerable constraint to always require contention avoidance.

09875436 "060604  
T09090" 9972960

In SIMD it is necessary for all PPs 100-103 to wait while contention is resolved. To achieve this a "SIMD pause" signal 3007 is routed between PPs 100-103, which can be activated by any PP 100-103 until their contention is resolved. Similarly in MIMD when executing synchronized code all synchronized PPs must wait until contention is resolved. This is signalled via sync signals 40.

The crossbar accesses are completed as soon as global ports 3005 and local ports 3006 are granted ownership of the RAM(s) they are attempting to access. In the case of stores they complete to memory 10 as soon as they are able. In the case of loads, if the PP is unable to resume execution immediately (because contention is continuing on the other port, or the SIMD pause signal 3007 is still active, or synchronized MIMD PPs are waiting for another PP, or a cache-miss has occurred) then the load(s) complete into holding latches 3018 and 3019 until execution is recommenced. This is because the data unit operation is also being held and its source data (i.e., a data register 3300) cannot be overwritten by a store. Similarly if a load and store are accessing the same data register and the store is delayed by contention, then the load data must be held temporarily in latch 3018 or 3019.

## 25 Data Unit

The logic within data unit 3000 works entirely during the execute pipeline stage. All of its operations use either registers only, or an immediate and registers. Indirect (memory) operands are not supported. Data transfers to and from memory are thus specifically coded as stores and loads.

A block diagram of data unit 3000 is given in FIGURE 33. The major components of the unit consist of 8 Data registers 3300, 1 full barrel shifter 3301, a 32-bit ALU 3302, a single-cycle 16x16 multiplier 3304, special

5 two devices when certain instructions are executed.

10 ALU 3302 and multiplier 3304 at the same time as two transfers to/from memory are occurring.

15 unsigned arithmetic.

20 different in that they support individual byte multiplexing. This is so that the "merge multiple (MRGM)" instruction can operate. This instruction uses the 4, 2 or 1 least-significant bits of the MFLAGS register to multiplex the individual bytes of each source with all zero  
25 bytes, so that what is passed into the ALU on one input is src1 bytes and 00h bytes intermixed according the M FLAGS. The opposite mix of 00h bytes and src2 is passed into the other ALU input. ALU 3302 can then do an ADD or an OR to produce a result which has some bytes from src1 and the  
30 others from src2. This is very useful for performing saturation, color expansion and compression, min and max, transparency and masking.

35 shift and add, or shift and subtract operations using a



5

10

15

20

25

30

35

32 bits each. The assigned codings are 000 - 2 bits, 001 - 4 bits, 010 - 8 bits, 011 - 16 bits, 100 - 32 bits. In the current implementation, however, the only permitted values are 8, 16 and 32 bits. These bit values control the operation of the ADDM, SUBM, MRGM and CMPM instructions.

#### Merge Multiple Instruction

Figure 44 shows some complex operations that can be performed by the combination of the splittable ALU instructions that set the MFLAGS register with the Merge Multiple (MRGM) instruction utilizing the multiplexer hardware of figure 33. The examples show only the data manipulation part of what would generally be a loop involving many of these operations.

In the add with saturate example of figure 44, the ADDM instruction does 4 8-bit adds in parallel and sets the MFLAG register according to whether a carry out (signalling an overflow) occurs between each 8-bit add. The 8-bit addition of Hex 67 to Hex EF and Hex CD to Hex 45 both cause a carry out of an 8-bit value which causes MFLAG bits 0 and 1 to get set (note only the 4 least significant bits of the MFLAG register will be significant to the MRGM instruction) resulting in the MFLAG register being set to "3". With D3 previously set to Hex FFFFFFFF, the MFLAG register values are used to select between the result of the previous operation contained in D2 or the saturation value of Hex "FF" stored in D3.

The Maximum function is obtained by doing a SUBM followed by using the same two registers with the MRGM instruction. The SUBM will set the bits of MFLAG register according to whether each 8-bits of a 32-bit value in one register is greater than the corresponding 8-bits in the other register as a result of 4 parallel 8-bit subtractions. As shown in the example, the MFLAG result of "5" (or binary "0101" for the 4 least significant bits)

5 final result stored in D2.

10 instruction performs 4 parallel 8-bit comparisons and sets the corresponding 4 MFLAG bits based on equal comparisons. In the example, only the third comparison from the right was "equal" signified by a "4" (binary "0100") in the MFLAG register. The MRGM instruction will then only use D0's  
15 values for the result except in the third 8-bits from the right.

20 moved into the MFLAG register. The MRGM instruction in this example simply selects between the 8-bit values in D0 and D1 according to the corresponding locations in the MFLAG register.

25 on whether or not the corresponding values match a specific  
color value. In this case the CMPM instruction's result in  
the MFLAG register is the result desired.

30 the destination. In the example the upper two 8-bit values  
of D0 will be copied to D1.

In the examples above 8-bit data values have been used by way of example. The number and size of the data values is not limited however to four eight-bit values.

Several important combinations of the arithmetic multiple instructions used with the merge instruction are shown. Many other combinations and useful operations are possible. It is significant that a large number of useful operations can be obtained by using the arithmetic multiple instructions that set the mask register and are followed by the merge instruction.

Two OPT bits in 3310 specify the type of shift that barrel shifter 3301 will perform during shift and add, and shift and subtract instructions. The codings are 00 - shift-right logical, 01 - shift-right arithmetic, 10 - shift-left logical, and 11 - rotate.

The AMOUNT bits in 3310 specify the number of bits of shift or rotate of the type indicated by the OPT bits, and occurring when shift and add, or shift and subtract instructions are executed.

### Appendix

The Appendix details each available instruction of the PPs 100-103. Dots (.) represent operation codes that can be assigned as desired. Some of these instructions have already been explained in the earlier text.

The order of instruction presentation is:

1. Data unit instructions (with or without parallel transfers) and single operation instructions (i.e., no parallel operations).

2. The transfers that can occur in parallel with data unit operations.

### 30 Transfer Processor

Transfer processor 11 is the interface between system memory 10 and the external world. In particular, it is responsible for all accesses to external memory 15.

Transfer processor 11, shown in detail in FIGURE 57, mainly performs block transfers between one area of memory

09875136-050603

and another. The "source" and "destination" memory may be on- or off-chip and data transfer is via bus 5700 and FIFO buffer memory 5701. On-chip memory includes: crossbar data memory 10, PP's instruction caches 10, master processor instruction cache 14, and master processor data cache 13 (shown in FIGURES 1 and 2). Data memories 10 and data cache 13 can be both read and written. The instruction caches 14 are only written.

All operations involving the caches are requested automatically by the logic associated with the caches. In this case the amount of data moved will be the cache "line" size, and the data will be moved between external memory 15 specified by the appropriate segment register and a segment of the cache.

Transfers involving crossbar data memories 10 are performed in response to "packet requests" from parallel processors 100-103 or master processor 12 and are accomplished via bus 5707. The packet request specifies the transfer in terms of a number of parameters including the amount of data to be moved and the source and destination addresses.

#### Block Transfers

A packet request specifies a generalized block transfer from one area of memory to another. Both source address generator 5704, and destination address generator 5705 are described in the same way. A "block" may be a simple contiguous linear sequence of data items (bytes, half-words, words or long-words) or may consist of a number of such regions. The addressing mechanism allows an "array" of up to 3 dimensions to be specified. This allows a number of two dimensional patches to be manipulated by a single packet request.

100-103-060601

Data items along the innermost dimension are always one unit apart. The distance between items of higher dimensions is arbitrary.

The counts of each dimension are the same for both  
5 source and destination arrays.

FIGURE 45 is an example of a complex type of block that can be specified in a single packet request. It shows a block consisting of two groups of three lines each consisting of 512 adjacent pixels. This might be needed  
10 for example if two PPs were going to perform a 3 x 3 convolution, each working on one of the groups of lines.

The block is specified in terms of the following parameters as shown in FIGURE 45:

15	Run length	Number of contiguous items e.g. 512 pixels.
	Level 2 Count	Number of "lines" in a group, e.g., 3
20	Level 3 Count	Number of "groups" in a "block" e.g., 2
	Start Address	Linear address of the start of the block, e.g., address of pixel indicated as "SA".
25	Level 2 Step	Distance between first level groups, e.g., difference of the addresses of pixels "B" and "A".
30	Level 3 Step	Distance between second level groups, e.g., difference of the addresses of pixels "D" and "C".

#### VRAM Auxiliary

35 The manner in which a video RAM would be used in conjunction with the multi-processor is described with respect to FIGURE 58 where the CCD input from the video camera or other video signal input would be clocked by AD converter 5802 into shift register 5801. Data can be  
40 shifted in or out of shift register 5801 into random memory matrix 5800 which in this case is the entire memory 15

shown in FIGURE 1. The S clock input is used to control the shifting of the information in or out shift register 5801. Data out of the random memory matrix 5800 is controlled by the parallel processors in the manner previously discussed such that the information can be used in parallel or in serial to do image processing or image control or figure identification or to clean the specks from paper or other copies. The ISP accesses the data in the video RAM via port 21 in FIGURE 58. The purpose of the shift register interaction with the random memory matrix is so that information can come asynchronously from the outside and be loaded into random memory matrix without regard to the processor operational speed. At that point the transfer processor then begins the transfer of information in the manner previously discussed. The input information would typically include NTSC standards which would include the horizontal sync and blanking and vertical refresh signals, which could be used as timing signals to control the loading or unloading of information from random memory matrix 5800.

The parallel processors can do many things with the data in random memory matrix 5800. Some of these can be processed at the same time. For example, color information can be separated for later processing or for distribution in accordance with the intelligence of the data, as previously discussed, or the information content of the received data can be manipulated as discussed previously with respect to FIGURE 11.

### 30 Operational Relationships

The number of controllers and data paths, and how they are configured with memory can be used to help classify architectures with respect to MIMD and SIMD. In simplest form a "processor" consists of one or more controllers and one or more data paths.

FIGURE 59 shows a typical MIMD configuration of four separate processing elements (5901, 5911, 5921, and 5931) connected to instruction (5904, 5914, 5924, and 5934) and data (5907, 5917, 5927, and 5937) memories. Note while the instruction and data memories are shown separately, they may actually be the same physical memory. Each processing element consists of two major blocks, the controller (5902, 5912, 5922, 5932) and data path (5905, 5915, 5925, 5935). The instruction memories provide control instructions to their respective controllers via instruction buses (5903, 5913, 5923, 5933). The data memories are accessed under control of the respective controller and go to the data paths via the data buses (5906, 5916, 5926, 5936). In some instances the instruction bus and data bus may in fact be the same physical bus, or the bus may actually be a set of buses configured in a crossbar arrangement. The controller controls the data path with a set of control signals (5908, 5918, 5928, 5938).

In the MIMD configuration of FIGURE 59, each processor can be executing completely independent instructions on either distributed or shared data.

FIGURE 60 shows a general SIMD configuration with a single controller 6002 and instruction memory 6004. Instructions pass to the controller via bus 6003. The single controller generates a single set of control signals 6000 that drive multiple data paths (6010, 6020, 6030, and 6040). Each data path is shown connected to its own memory (6012, 6022, 6032, 6042) via buses (6011, 6021, 6031, 6041). While for simplicity each data path is shown having a single way of connecting to the data memories, there may in fact be various ways in which the data paths and data memories can be connected such as via a crossbar arrangement or via a sequential passing of data as shown in FIGURE 8.



In the SIMD configuration of FIGURE 60, a single instruction stream is used to control multiple data paths. In the general SIMD case, such as shown in FIGURE 60, there is only one controller for the multiple data paths.

5       FIGURE 61 shows an embodiment of the system which is the subject of this invention, where the system is configured to behave in a MIMD mode. Via the crossbar 20, each parallel processor (100, 101, 102, or 103) can each use a memory within the memory space 10 as its instruction  
10 memory. The controller 3002 of each parallel processor thus can get its own different instruction stream. The synchronization signals in bus 40 are ignored by each parallel processor that is configured to be in the MIMD mode of operation. Since each controller can control via  
15 control signals 3112 a different data path 3100 and each data path can have access to a different memory via the crossbar, the system can operate in a MIMD mode.

FIGURE 62 shows the same hardware of FIGURE 61, however, the parallel processors have been configured in a  
20 SIMD mode. In this mode, a single instruction memory is connected to all processors as described in the discussion related to FIGURE 28. With each of the SIMD organized parallel processors receiving the same instruction, each controller will issue generally the same control signals.  
25 For example, there may be differences in control signals due to data dependencies which must be taken account of. The synchronization signals in bus 40 serve two purposes: first they are used to get the parallel processors all started on the same instruction when transitioning from  
30 MIMD to SIMD operation, and second once started in SIMD operation they keep the parallel processor from getting out of step due to events that may not affect all processors equally (for example if two processors access the same memory, the conflict resolution logic will allow one of the  
35 processors to access the memory before the other one).

09375136 "060604  
T09090 SETS480

5

10

20

### Preferred Embodiment Features

30

35

one or more addressable memories with  $m$  memory sources each having a unique addressable space, where  $m$  is greater than  $n$  and having a switch matrix connected to the memories and connected to the processors and with circuitry for  
5 selectively and concurrently enabling the switch matrix on a processor cycle by cycle basis for interconnecting any of the processors with any of the memories for the interchange between the memories and the connected processors of instruction sets from one or more addressable memory spaces  
10 and data from other addressable memory spaces.

A processing system is shown with a plurality of processors, arranged to operate independent from each other from instructions executed on a cycle-by-cycle basis, with the system having a plurality of memories and circuitry for  
15 interconnecting any of the processors and any of the memories and including circuitry for interconnecting any of the processors and any of the memories and including circuitry for arranging a group of the processors into the SIMD operating mode where all of the processors of the  
20 group operate from the same instruction and circuitry operable on a processor cycle-by-cycle basis for changing at least some of the processors from operation in the SIMD operating mode to operation in the MIMD operational mode where each processor of the MIMD group operates from  
25 separate instructions provided by separate instruction memories.

An image processing system is shown with  $n$  processors, each processor operable from instruction streams provided from a memory source for controlling a number of different  
30 processes, which processes rely on the movement of data from  $m$  addressable memories each having a unique addressable space, and wherein  $m$  is greater than  $n$  and with a switch matrix connected to the memories and connected to the processors and including circuitry for selectively and  
35 concurrently interconnecting any of the processors with any

of the memories so that the processors can function in a plurality of operational modes, each mode having particular processor memory relationships; and including an interprocessor communication bus for transmitting signals  
 5 from any processor to any other selected processor for effecting said operational mode changes.

A multi-processing system is further shown comprising n processors, each processor operable from an instruction stream provided from a memory source for controlling a  
 10 process, said process relying on the movement of data to or from m addressable memories; each memory source having an addressable space and a switch matrix having links connected to the memories and connected to the processors; and including circuitry for splitting at least one of the  
 15 links of the switch matrix for selectively and concurrently interconnecting any of the processors with any of the memories for the interchange between the memories and the connected processors of instruction streams from one or more memory addressable spaces and data from other  
 20 addressable memory spaces.

A processing system is shown having a plurality of processors, each processor capable of executing its own instruction stream with control circuitry associated with each of the processors for establishing which of the  
 25 processors are to be synchronized therewith and with instruction responsive circuitry associated with each processor for determining the boundary of instructions which are to be synchronized with the other synchronized processors and for setting a flag between such boundaries;  
 30 and including circuitry in each processor for establishing a ready to execute mode; and control logic associated with each processor for inhibiting the execution of any instruction in the processor's instruction stream while each flag is set in the processor until all of the other  
 35 processors established by the processor as being

synchronized with the processor are in a ready to execute mode.

A multi-processing system is shown with  $m$  memories, each memory having a unique addressable space, with the  
 5 total addressable space of the  $m$  memories defined by a single address word having  $n$  bits; and a memory address generation circuit for controlling access to addressable locations with the  $m$  memories according to the value of the bits of said address word; and with addition circuitry  
 10 having carryover signals between bits for accepting an index value to be added to an existing address word to specify a next address location; and with circuitry operative for diverting the carryover signals from certain bits of said word which would normally be destined to  
 15 toggle a next adjacent memory address word bit so that said carryover signal instead toggles a remote bit of the memory address word.

A circuit for indicating the number of "ones" in a binary string, the circuit having an AND gate having first  
 20 and second inputs and an output; an XOR gate having first and second inputs and an output, the first input thereof connected to the first input of the AND gate, the second input connected to the second input of the AND gate; and where the second inputs of the AND and XOR gates receive  
 25 one bit of the binary string and the output of XOR gate produces an output binary number representative of the number of "ones" in the bit of the binary string.

A multi-processing system is shown with  $n$  processors operable from instruction streams provided from a memory  
 30 source for controlling a number of different processes, said processes relying on the movement of data from one or more addressable memories; and with  $m$  memory sources, each having a unique addressable space, some of the memories adapted to share instruction streams for the processors and  
 35 the others of the memories adapted to store data for the

processors; and with a switch matrix for establishing communication links between the processors and the memories, the switch matrix arranged with certain links providing dedicated communication between a particular processor and a particular one of the memories containing the instruction streams; and with circuitry for rearranging certain matrix links for providing data access to memories previously used for instructions, and circuitry concurrently operative with the rearranging circuitry for connecting all of the processors to a particular one of the certain links so that instructions from the instruction memory associated with the certain link are communicated to all of the system processors.

An imaging system having an image input, each image having a plurality of pixels, each pixel capable of having a plurality of data bits associated therewith; a memory; an image bus for transporting pixels from each image at the input to the memory; and circuitry for interpreting received images in accordance with parameters stored in the memory, the interpreting resulting from the parameters being applied to the pixels of each received image.

A switch matrix is arranged for interconnecting a plurality of first ports with a plurality of second ports, the switch matrix having: a plurality of vertical buses, each bus associated with a particular one of the first ports; and a plurality of individually operable crosspoints; and a plurality of horizontal buses connected to the second ports for connecting, via enabled ones of the cross points, one of the first ports and any one of the second ports and including circuitry at each crosspoint, associated with each vertical bus for handling contention between competing ones of the second ports for connection to said vertical bus.

0987535-050604

Summary

Although the present invention has been described with respect to a specific preferred embodiment thereof, various changes and modifications may be suggested by one skilled in the art, and it is intended that the present invention encompass such changes and modifications as fall within the scope of the appended claims. Also, it should be understood that while emphasis has been placed on image processing the system described herein can as well be used for graphics, signalling processing, speech, sonar, radar and other high density real time processing. High definition TV and computing systems are a natural for this architecture.

T.03030 "92.7.52.860